RTC
The magazine of record for the embedded computing industry

Also in this issue

October 2005 www.rtcmagazine.com

VME's Next-Generation READY for the Real World

RTC Interviews Curtiss-Wright's Tom Quinly

An RTC Group Publication

Who says you can't plug a PC into a VME chassis?

SBS knows how to get buses to talk to each other.
Control your VME chassis with our PCI card.

INSTALL TWO BOARDS. Connect them with a fiber optic cable. Fire up the system. And next thing you know you'll be sitting comfortably in front of your PC making your VME system jump through its usual hoops. It's almost like plugging a PC into your VME system. Wait a minute, it's exactly like that.

With this adapter, off-the-shelf software, any PC you like and ordinary peripherals, you can create a machine-human interface at truly affordable price and performance levels. And our drivers make the integration process painless. It's a simple, economical way to bring all the things you love about your desktop directly to your embedded applications.

VME64 to PCI
bus-to-bus adapter with DMA

This new SBS adapter features memory and I/O mapping, controller mode DMA, interrupts and portable drivers. It directly connects the two buses, which allows the two systems to operate as one. The result is a real-time and deterministic system with low latency and low CPU overhead. The point-to-point connectivity is at the hardware level so no protocol is required and the adapter never detracts from normal system performance. Pretty cool, huh?

SBS Technologies®

SBS knows. Find our VME bus adapters at **www.sbs.com/busadapter** or call **800.SBS.EMBEDDED**
Order your free VME Reference Book at **www.sbs.com/vmebook**

Cool. Calm. Collected.

The Themis Slice™ architecture can handle whatever you or the outside world throw at it.

Cool.

Innovative and cost-effective cooling system can tame the hottest of processors in four 100 Watt sockets that support eight processor cores. Air and liquid heat exchanger options.

Calm.

Standardized mechanical footprint designed to survive in high shock and vibration environments - up to 40G shock and 1 grms vibration.

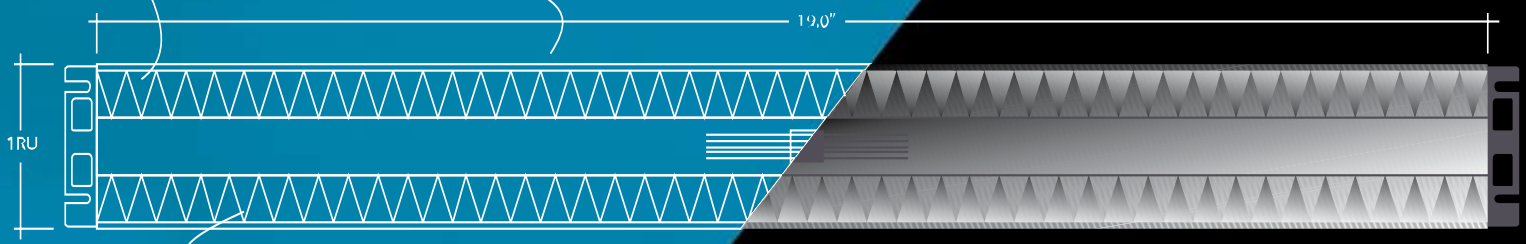

Collected.

Scalable - up to 16 processor cores. High-density, processor-independent architecture implemented with SPARC®, IBM®, PowerPC and x86 technology. Solaris™, Windows®, and Linux® operating systems. All networked in a single system. All managed in real-time.

Themis' Slice technology is a breakthrough in lifecycle management for mission-critical computing.

By functionally disaggregating commercial computing resources and housing them in a standardized footprint, purpose-built enclosure, Slice provides resilience with superior thermal and kinetic management. This open and modular design allows for spiral technology refresh, extending computing infrastructure investments for complete lifecycle management.

Themis' real-time resource manager integrates all computing resources, under dynamic policy control, to optimize application quality of service.

Themis Slice is transformational technology offering the most robust and efficient platform available. No other platform will enable the 21st century war fighter as effectively.

www.themis.com (510) 252-0870

THEMIS

Transformational.

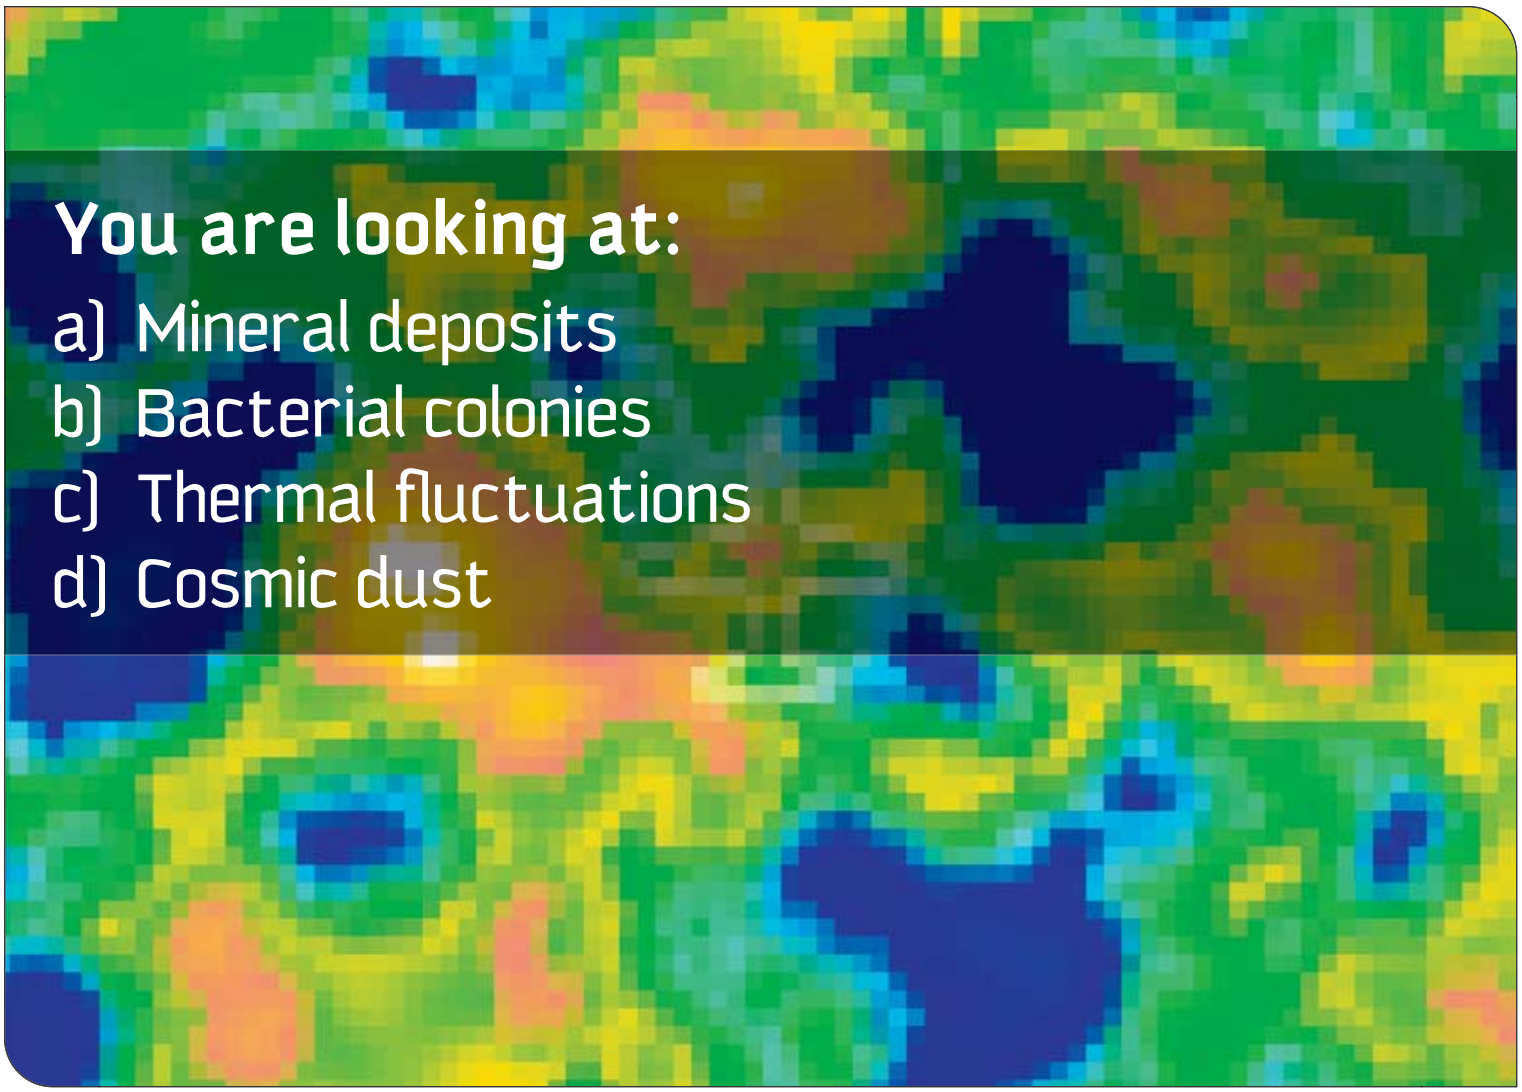

SBS knows Graphics. Visit sbs.com/g2 for the answer to the above, and how the G2 can solve your graphics challenges.

NO MATTER WHAT YOUR APPLICATION, our new graphics PMC has the flexibility and processing power you need. It can handle multiple inputs and push multiple displays, thanks to the ATI® Mobility™ Radeon® 9000 (M9) chip. It accepts LVDS, RGB and Composite video inputs in a number of different formats and frame rates, and drives dual displays in multiple formats. This family of graphics boards comes in 3 models and each model is available in 3 environmental grades: Commercial, Extended Temperature, and a rugged Conduction Cooled version.

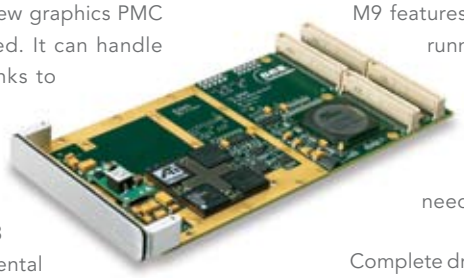

G2 Plus Graphics PMC Module

The M9's internal processors provide separate concurrent pipeline engines for 2D processing, 3D processing, output processing, and video capture processing. The M9 features 64 MBytes of 128-bit wide DDR video memory running at 200 MHz. Combined, this provides graphics power approaching that of a high-end desktop workstation, all on a PMC module. SBS routes the video capture streams through a configurable FPGA to provide the design flexibility designers need for today's systems.

Complete driver software support and demonstration software from ALT Software for SBS' customers, ensures trouble-free operation combined with optimized performance in Microsoft® Windows®, Linux®, VxWorks® and INTEGRITY®. ALT Software and SBS can even provide DO-178B support. Put all that together, and what you're really looking at is a whole new level of embedded graphics capability.

SBS Technologies.

SBS knows.
See the G2 Plus in action at **www.sbs.com/g2** or for more information call **800.SBS.EMBEDDED**

Departments

RTC

www.rtcmagazine.com

Features

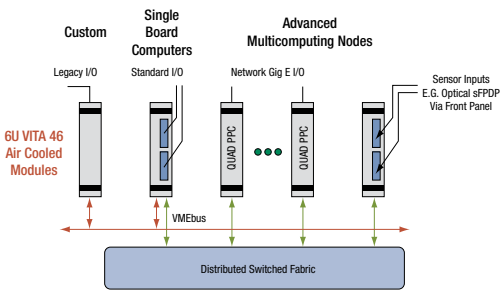

Radar processing—VITA 46 enables the use of next-generation switched fabrics such as Advanced Switching Interconnect, Serial RapidIO, etc., in either centralized or distributed architectures. • **Pg. 36**

The CPC6314 from Performance Technologies is a 325W DC power supply that combines high power capacity and industry standard IPMI management capability with monitoring and reporting. • **Pg. 50**

Rugged PMC Modules Offer 3D Graphics Processing • **Pg. 70**

Publisher
PRESIDENT
John Reardon, johnr@rtcgroup.com

VICE PRESIDENT, European Operations
Zoltan Hunor, zoltanh@rtcgroup.com

EDITORIAL DIRECTOR/ASSOCIATE PUBLISHER
Warren Andrews, warrena@rtcgroup.com

Editorial
EDITOR-IN-CHIEF
Tom Williams, tomw@rtcgroup.com

SENIOR EDITOR
Ann Thryft, annt@rtcgroup.com

MANAGING EDITOR
Marina Tringali, marinat@rtcgroup.com

STAFF EDITOR
Craig Choisser, craigc@rtcgroup.com

COPY EDITOR
Rochelle Cohn

Art/Production
CREATIVE DIRECTOR
Jenna Hazlett, jennah@rtcgroup.com

GRAPHIC DESIGNER
Melissa Gaeta, melissag@rtcgroup.com

PRODUCTION DESIGNER
Kirsten Wyatt, kirstenw@rtcgroup.com

DIRECTOR OF WEB DEVELOPMENT
Jan Garcia, jang@rtcgroup.com

WEB DEVELOPER
Marke Hallowell, markeh@rtcgroup.com

Advertising/Web Advertising
CALIFORNIA COASTAL ADVERTISING MANAGER
Diana Duke, dianad@rtcgroup.com
(949)226-2011

WESTERN REGIONAL ADVERTISING MANAGER
Lea Ramirez, lear@rtcgroup.com
(949)226-2026

EASTERN REGIONAL ADVERTISING MANAGER
Nancy Vanderslice, nancyv@rtcgroup.com
(978)443-2402

EMEA SALES MANAGER
Anna Brown, annab@rtcgroup.com
(949)226-2025

ASIAN ADVERTISING MANAGER
John Koon, johnk@rtcgroup.com
(858)755-4999

Billing
Maggie McAuley, maggiem@rtcgroup.com
(949)226-2024

To Contact RTC magazine:

HOME OFFICE
The RTC Group, 905 Calle Amanecer, Suite 250, San Clemente, CA 92673
Phone: (949) 226-2000 Fax: (949) 226-2050, www.rtcgroup.com

EASTERN SALES OFFICE
The RTC Group, 96 Dudley Road, Sudbury, MA 01776
Phone: (978) 443-2402 Fax: (978) 443-4844

EUROPEAN OFFICE
The RTC Group Europe, PO BOX 2114, NL 5300 CC Zaltbommel, Netherlands, Phone: 31-418-577700 Fax: 31-418-577720, www.rtcgroup.com

Editorial Office
Warren Andrews, Editorial Director/Associate Publisher
39 Southport Cove, Bonita, FL 34134
Phone: (239) 992-4537 Fax: (239) 992-2396

Tom Williams, Editor-in-Chief
245-M Mt. Hermon Rd., PMB#F, Scotts Valley, CA 95066
Phone: (831) 335-1509 Fax: (408) 904-7214

Ann Thryft, Senior Editor
15520 Big Basin Way, Boulder Creek, CA 95006
Phone: (831) 338-8228

Published by The RTC Group

Looking for a solution?

We've got it!

Electronic Solutions continues to reach new heights in standard off-the shelf chassis products supporting the Open Bus Architectures for CompactPCI, VME, VME64Extensions, PXI/VXI, AdvancedTCA, and other Switched Fabric Architectures. Our system level components and turnkey system solutions incorporate high performance interconnect technologies for today's fast moving industries: telecommunications, networking, commercial, instrumentation, medical and military.

At Electronic Solutions our focus is on the optimum solution in support of your unique system packaging requirements. We can make modified or custom solutions to your requirements.

Electronic Solutions

a business unit of apw Ltd.

ARCHITECTURES

ATCA | cPCI | VME | VME64X | VXI / PXI

For more information: **sales.elsol@apw.com** or visit **www.electronicsolutions.com**

VMEbus Pentium M

Low-Power/High-Performance

As Rugged As You Need It!

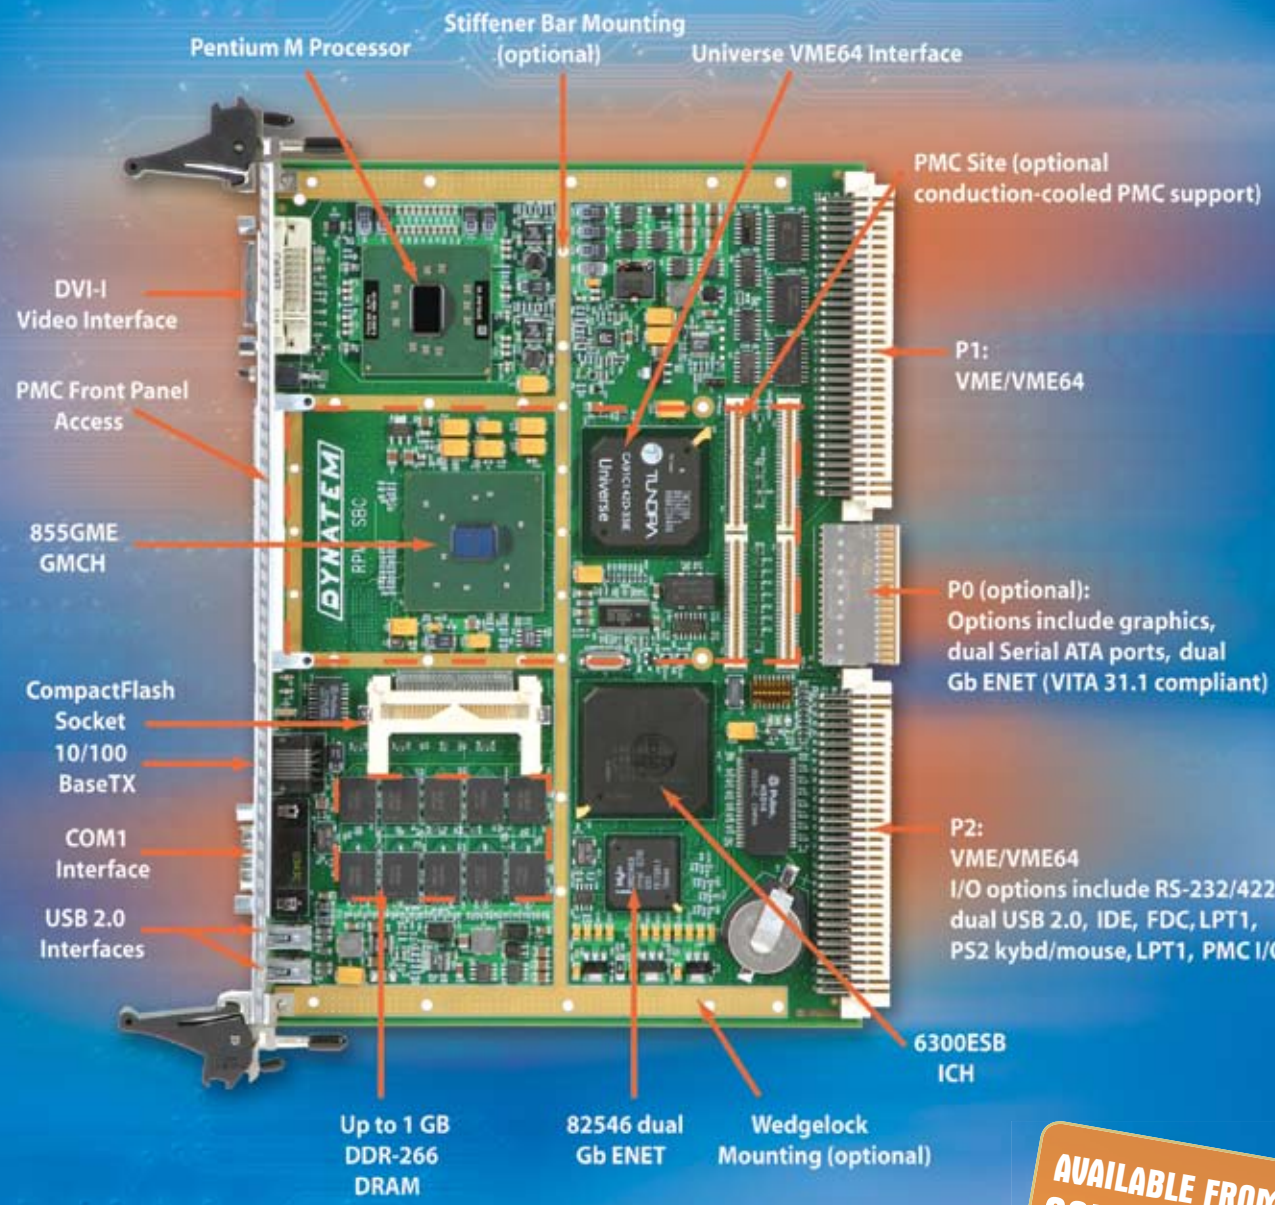

AVAILABLE FROM COMMERCIAL TO FULL MIL

DPM FEATURES

- NT, XP, Linux, VxWorks and Solaris all run on-board and there's PXE support for diskless booting in secure systems
- Two Gb Ethernet ports in compliance with VITA 31.1 for backplane fabric switching
- USB 2.0, DVI-I Graphics, Serial ATA, and PCI-X are integrated into the Chipset
- Rugged, extended temperature versions are available with wedgelocks and conduction cooling in compliance with IEEE 1101.2. Speed-step capable.

DYNATEM

www.dynatem.com 800.543.3830 949.855.3235 sales@dynatem.com

Dynatem, Inc., 23263 Madero, Suite C, Mission Viejo, CA 92691

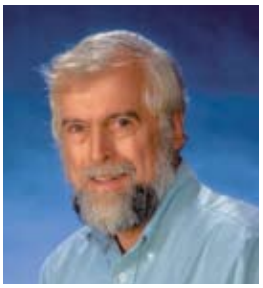

The Little Form-Factor That Could

by Tom Williams, Editor-in-Chief

In an industry roiling with turbulent acquisitions, savage takeovers, ruthless patent disputes and cutthroat competition, and additionally beset by commoditization, how is it that there can be a group of some 50 moderate-sized companies, almost all of them privately held, that are fiercely competitive yet seem to all be fairly successful? At least they appear to be successful because most of them have been around for a while. I'm talking about the PC/104 sector. While there are some public companies that manufacture PC/104 modules, these often do it as a part of a larger product offering. Those companies whose main business is PC/104 are, as far as I know, all in private hands.

If you're selling CompactPCI into the telecom market, you have a whole lot of challenges. In addition to competition, there is the need to make volume sales, but to do that you also have to offer something with unique features to suit a customer's needs. Since different customers have different needs, that means a range of focused products that need to sell in high volumes to make a profit. That means you must not only have the insight and engineering capability to design the needed products on time in the face of the competition, but you must also be able to supply them in the volume needed. So manufacturing capacity may reach its limits and lead to outsourcing.

The other danger is the flip side of the first. That is that a product is so general-purpose that it can be manufactured in quantity by various companies and be differentiated by a combination of software or add-on modules such as PMC cards. This commoditization drives down the margins and again favors the bigger companies driving the smaller ones to failure or, if they are lucky, acquisition. Telecom is one of those application areas that as an industry is happy to let smaller companies do its R&D for it and then move to contract manufacturing when decisions to purchase in volume ramp up.

VME has avoided the commoditization trap by addressing markets that have inherently specialized needs such as the military with its ruggedization requirements, etc., or specialized niche applications like radar and software-defined radio. Here the volumes are relatively low and the margins correspondingly high.

PC/104 has found a sweet spot. Its manufacturers can produce relatively large numbers of small, low-cost modules that can be mixed and matched by a wide variety of customers to meet a vast array of application requirements. The goal was to spare the customer the time and effort of specialized engineering, and the result is that there is a large number of available CPU and I/O modules available from a large number of vendors. Having done the R&D on these, the PC/104 vendors can produce them in sufficient volume to feed a large number of applications with different needs. It's sort of like a niche market with a large number of sub-niches. In addition, these companies are quite happily able to sell 50 of this and 75 of that (not that they wouldn't appreciate larger orders) and keep acceptable margins.

With the advent of EPIC we have seen the introduction of a backward-compatible form-factor just large enough to accommodate the more powerful CPUs and additional standard and custom I/O. Now, with the addition of EPIC Express, we have the path to the most popular serial switched fabric, which is about to move to 5 Mbit/s speeds. In addition, the EPIC form-factor is starting to appear in versions that do not strictly adhere to the spec except with respect to the form-factor. These allow more specialized customization of I/O with a minimum of pain to the customer

As the I/O bandwidth capabilities increase, we will probably see a move to more powerful CPUs in order to keep up with it, and that could bring along some issues of power consumption and cooling. That remains to be seen. Still, the PC/104 family is set to move into the future while maintaining backward compatibility and broad applicability for customers whose main goal is to add value beyond the hardware configuration. Those who wish to do custom engineering in this space are, of course, free to do so. However, the cost and performance levels are such that a solution need not be absolutely optimal in order to be of good value in terms of cost, space, power and time-to-market.

Additionally, I've been struck by the fact that many of the principles of these companies, some of whom I know personally, appear not to be suffering from stress disorders, have good senses of humor, seem genuinely happy at what they are doing and relatively prosperous. In the grand scheme of things, who could ask for more?

Real VME Solutions.

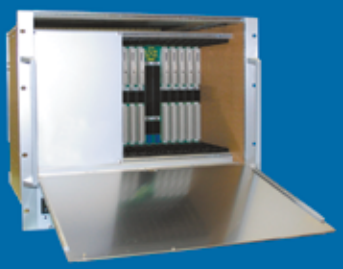

VXS

VME

VME64X

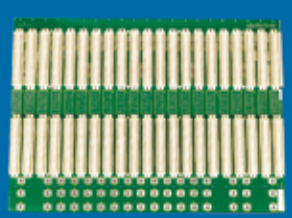

BACKPLANES

Real Experience. *Real* Results.

You've heard it before. Others have claimed to have the solution. But there's a reason Elma is *the* supplier of choice for VME enclosures in the U.S. We get the job done. We've been involved with VME since its inception, and we continue to actively participate in the development of cutting edge technologies that refine platform configurations. A few of our latest projects include VME64X, VITA 31 (GigE), VITA 41 (VXS; switched serial standard) and VITA 46. The Elma commitment to customer-oriented services, and better product performance means a real payoff for you. Call us today. Elma delivers on its word.

Enclosures & Components | Backplanes | System Platforms | Switches, Knobs & LEDs | Cabinets

ELMA
Your Solution Partner

Elma Electronic Inc.
Phone: 510.656.3400 Fax: 510.656.3783 E-Mail: sales@elma.com Web: elma.com

INDUSTRY INSIDER

October 2005

Rumblings on the Horizon for AMC?

There appear to be a number of issues affecting the recently introduced Advanced Mezzanine Card (AMC) standard. A number of companies, including Intel, are calling for a large number of major and minor revisions to the specification. At the root of the question appears to be the notion that although AMC was specified to work well in networking and telecom environments—often in NEBS-compliant buildings—people are starting to look at it as a possible solution in military and industrial control applications as well. This has raised doubts as to how well AMC will hold up under the additional stresses imposed by these environments.

Among the issues are shock and vibration, which impact the design of the connector as well as the front panel latches and other mechanicals. The dilemma that faces the industry appears to be: will these issues really prevent its use in the expanded universe of applications, and is it compelling enough to force a total re-work of the specification? On the other hand, since AMC is poised to enter the market in a large number of already identified roles, what will be the cost of delaying that entry in order to re-do the spec? *RTC* will follow this subject as it evolves and invites qualified representatives in the industry to offer their points of view with an eye toward fostering a fair and productive discussion of issues that affect our industry.

Real-Time Innovations Acquires Technology for Real-Time Data Management

Real-Time Innovations has announced the acquisition of 4TEC BV's SkyBoard intellectual property (IP) and development team. SkyBoard distributes databases across multiple nodes, enabling performance-critical applications to store and access data from distributed servers at

extremely fast rates. SkyBoard provides infrastructure software for real-time information management systems supporting performance-critical applications such as public safety, global communications, aerospace, financials, manufacturing and transportation.

With the SkyBoard acquisition, RTI now offers infrastructure software that acts as a real-time information management system for dynamic, heterogeneous networks. The system supports topologies that continuously change as devices join and leave the network. Improved performance over existing systems makes it possible to develop increasingly demanding applications that encompass an ever-growing number of devices in widely dispersed networks.

Under the Object Management Group's (OMG) Data Distribution Service (DDS, V1.0 standard), developers specify various quality-of-service (QoS) parameters such as reliability versus best-effort delivery, expected rates of publication and length of time data is valid throughout the network. Modules sending data become publishers and are only concerned with the specific data type they communicate. Modules needing specific information are subscribers and only need to know about the particular data they wish to receive.

SiliconSystems Enters Long-Term Arrangement with Samsung

SiliconSystems, a manufacturer of industrial-grade solid-state storage technology, has announced that they have agreed to work with Samsung Semiconductor over the next several years to aggressively drive solid-state technology use among system OEMs. Under this arrangement, Samsung will increase its supply of SLC NAND flash for use in SiliconSystems' SiliconDrive products to meet the increasing demand of system OEMs for SLC-based storage solutions. The enterprise OEM storage market is expected to approach $2.6 billion by 2008, according to Web-Feet Research, and is growing at an annual growth rate of 48.1 percent.

SiliconSystems' SiliconDrive provides increased performance and industrial-grade reliability for embedded and industrial applications. SiliconDrives offer endurance and reliability, and are designed specifically to meet the rigorous demands of data-rich and mission-critical applications in the system OEM market. With a 6-bit ECC algorithm, SiliconDrives feature wear-leveling technology and an interface between the solid-state storage array and host system. To eliminate drive corruption and ensure data integrity and storage availability, SiliconDrives integrate SiliconSystems' PowerArmor technology. Additionally, SiliconDrives are enhanced by integrating application-specific technologies such as data sanitization and purge, secure write protection and faster transfer speeds for small data file transactions.

Program Enables Carriers to Begin Selling IP Services

VoIP, Inc. has announced the launch of its VoiceOne Carrier Direct Program. The new program enables carriers for IP, providing them immediate availability to these services with no CAPEX or other requirements. The Carrier Direct Program provides carriers with media gateways at no charge, and installs and connects them to the VoiceOne network. This enables carriers to immediately begin selling IP-based services, such as hosted IP centrex, broadband voice, IP origination/termination, 800 service, CNAM, 911 services, etc. over the VoiceOne network.

The program also enables carriers to increase their level of service offerings and recognize an immediate ROI, utilizing VoIP's technical support as they gain entry into this market. According to VoIP, one advantage of this service for carriers is that it is integrated back into their existing TDM infrastructure with no changes and no CAPEX requirements. Monitoring and support are provided through the VoiceOne 24-hour Network Operations Center, which also alleviates additional operating expenses for carriers.

The VoiceOne network has its 22 points of presence in the U.S. and over 5,000 local access points reaching approximately 90 percent of the U.S. population. VoiceOne is a protocol-agnostic network, supporting virtually any protocol including SS7, PRI, H.323, SIP,

Event Calendar

10/11/05 (new date)
Real-Time & Embedded Computing Conference
Long Beach, CA
www.rtecc.com/longbeach

11/15/05
Real-Time & Embedded Computing Conference
Atlanta, GA
www.rtecc.com/atlanta

11/17/05
Real-Time & Embedded Computing Conference
Raleigh-Durham, NC
www.rtecc.com/raleighdurham

12/01/05
Real-Time & Embedded Computing Conference
Portland, OR
www.rtecc.com/portland

12/6/05
Real-Time & Embedded Computing Conference
Seattle, WA
www.rtecc.com/seattle

12/08/05
Real-Time & Embedded Computing Conference
Vancouver, BC
www.rtecc.com/vancouver

If your company produces any type of industry event, you can get your event listed by contacting sallyb@rtcgroup.com.

This is a FREE industry-wide listing.

MGCP, etc. and utilizes both circuit and packet switched technology.

Artesyn Partners with MontaVista on Carrier Grade Linux for ATCA and AMC

Artesyn Communication Products announced that it has partnered with MontaVista to offer MontaVista Linux Carrier Grade Edition (CGE) for its telecom blades and modules. Artesyn will offer bundled, certified MontaVista CGE solutions for its PICMG 2.16, AdvancedTCA, AdvancedMC, ProcessorPMC blades and modules, preinstalled in flash memory. Artesyn

blades and modules equipped with MontaVista CGE provide a modular open architecture platform for building scaleable, high-availability network infrastructure equipment.

MontaVista Linux CGE is an open and flexible development platform designed specifically to address the requirements of carrier grade class applications, with a strong focus on open standards and high-availability services. MontaVista CGE is the industry's most field-tested and proven Linux distribution. MontaVista CGE is available immediately from MontaVista for Artesyn's PICMG 2.16 Katana 750i, 752i, 3750 and 3752 blades, which feature PowerPC processing complexes, high-performance packet-switched backplane interfaces, PMC/PTMC mezzanine expansion and integrated IPMI system management. MontaVista CGE is also available for Artesyn's PowerQUICC-based Pm8560 protocol engine, an octal E1/T1 card optimized for SS7/SIGTRAN signaling. Artesyn will announce bundled, certified MontaVista CGE solutions for select PICMG 2.16, AdvancedTCA, AdvancedMC and ProcessorPMC blades and modules at a future date.

Green Hills Software and I-Logix Team to Offer IDE with UML and Multi Tools

Green Hills Software and I-Logix have entered a strategic partnership resulting in a single source, standards-based Integrated Development Environment (IDE) that addresses all phases of embedded systems development, from requirements specification through deployment. Under the agreement, Green Hills Software will distribute and support I-Logix' Model-Driven Development (MDD) environment, named Rhapsody, which is based on the Unified Modeling Language (UML) 2.0 standard, together with Green Hills Software's Multi and AdaMulti devaelopment environments as well as the POSIX-conformant Integrity RTOS. The products are integrated to enable a bi-directional workflow between modeling and implementation. The companies are also collaborating on future integrated features and capabilities, including integration with the Eclipse platform.

The unified solution developed by Green Hills Software and I-Logix generates the implementation in C, C++ or Ada source code directly from the UML model. The model is then automatically updated to reflect any changes made to the source code. The combination

of Rhapsody and Multi also helps ensure that a final product satisfies its design objectives by providing traceability between the source code and requirements. In addition, the Multi source code debugger as well as the graphical UML model can be used to set breakpoints and single-step through code.

VoX Communications Adds Enhanced 911 Service to VoIP Offering

VoX Communications, a provider of Voice-over-Internet Protocol (VoIP) services, has announced that it has selected Global Crossing's (Enhanced 911 (E911)) solution to support its VoIP services offering. VoX is a wholly owned subsidiary of eLEC Communications.

E911 is an emergency reporting system that reports the Automatic Number Identification (ANI) and the Automatic Location Identification (ALI) when connecting an emergency call to the assigned PSAP. This allows the PSAP to immediately identify the exact location of the caller's telephone and dispatch assistance at once, regardless of the caller's ability to communicate, when receiving a 911 call.

VoX Communications offers wholesale broadband voice, origination and termination services for cable, wireless and wireline operators and enhanced VoIP services to the small business and residential marketplace. VoX's offering includes Call Hold, Call Waiting, Caller ID, Call Transfer, Hunt Groups, Do Not Disturb, Call Forward, International Call Blocking, Call Return, Repeat Dialing/Redial, Extension Dialing, Anonymous Call Rejection and e-mail notification of voicemail. Add-on features include: Multibox Voicemail, Music on Hold, Corporate Conference calling, Reassign Phone, Find me/Follow me and Auto Attendant, among others.

4U 8-slot CompactPCI Subsystem

The cPCIS-6418U is a high-performance CompactPCI platform with a 64-bit, 8-slot 6U backplane in a 4U enclosure. The 8-slot backplane dedicates one system CPU slot and seven peripheral slots. Dual AC inlet and three redundant power supply units provide 500W+250W power redundancy. Featured with high availability and ample room for expansion, the cPCIS-6418U is compliant with the PICMG 2.5 computer telephony integration standards and is ideal for VoIP and network equipment applications.

For more info, go to:
www.adlinktech.com/products

2-CH 65MS/s 14-Bit Digitizer with SDRAM

ADLINK's PCI-9820 and PXI-9820 are 2-CH 65 MS/s, high-resolution PXI digitizer with 512MB SODIMM SDRAM memory. They features flexible input configurations including programmable input ranges and user selectable input impedance. With lots of on-board acquisition memory, these digitizers are not limited by the PCI 132MB/s bandwidth and can record the waveform for a long period of time. The PCI-9820 and PXI-9820 are ideal for high-speed waveform capturing, software radio applications, and signal digitizing applications which require ample memory.

For more info, go to:
www.adlinktech.com/products

8-CH High-Power Relay Outputs & 8-CH Isolated Digital Inputs Card

The PCI-7260 is a high-power relay output card for automation and machine control. It comes with eight channels switches 5A currents at 250VAC or 30VDC. The front panel connector and PCB are implemented for carrying high currents. The removable terminal block is convenient for easy wiring. The PCI-7260 also offers eight isolated digital input channels with debouncing circuits to eliminate signal interference.

For more info, go to:
www.adlinktech.com/products

An Entire Family of High-Speed, Low-Power Universal Intel® Pentium® M Processor Boards

The cPCI-6840 family combines Intel's latest high performance low power processor the Pentium® M with the 855GME MCH and two different I/O controller hubs. The Pentium® M has a highly efficient instruction execution unit and coupled with a large L2 cache, it can outperform other processors at the same clock speed. The 855GME has an integrated 250MHz 32-bit 3D/2D graphics engine capable of supporting simulations CRT and flat panel displays. The ICH packs a plethora of peripherals include Serial ATA 150, PATA, USB 2.0, watchdog timer and serial port.

For more info, go to:
www.adlinktech.com/products

A Variety of Intel® Pentium® M based Single Board Computers

3U PXI CPU Module
with VGA, GbE & HotSwappable CF

3U CompactPCI CPU Module
with Dual GbE & USB 2.0

Half-Size Low Power SBC
with Dual GbE, LVDS & CF

ETX Module with High-End Graphics
& Optional HDTV I/F

Call us toll-free at (866) 4-ADLINK
or email info@adlinktech.com

PXI Systems Alliance

PICMG

Intel® Communications Alliance
Associate Member
SILVER

ADLINK TECHNOLOGY INC.

www.adlinktech.com

EPIC + PCI Express = EPIC Express: The Industrial Computing Platform's Next Generation

PCI Express is the wave of the future for high-performance I/O in both desktop and industrial/embedded applications. Yet the stacking-style architecture employed with EPIC, PC/104 and EBX boards presents unique challenges, which EPIC Express solves. The logical evolution of the EPIC specification adds PCI Express expansion capability to this industrial computing platform.

by Robert A. Burckle
WinSystems

When EPIC—the Embedded Platform for Industrial Computing—was introduced, it called for a future upgrade to the specification to support a serial fabric. PCI Express technology has emerged as the next-generation I/O solution of choice in many of the computing and communications industries as it is now migrating from the desktop to embedded applications. The PCI Express architecture uses the familiar software and configuration interfaces of the conventional PCI bus architecture, but provides a new, high-performance physical interface while retaining software compatibility with the existing conventional PCI infrastructure. The reason for adding PCI Express was to provide a "bridge to the future" while maintaining legacy support for the vast number of PC/104 expansion modules available worldwide.

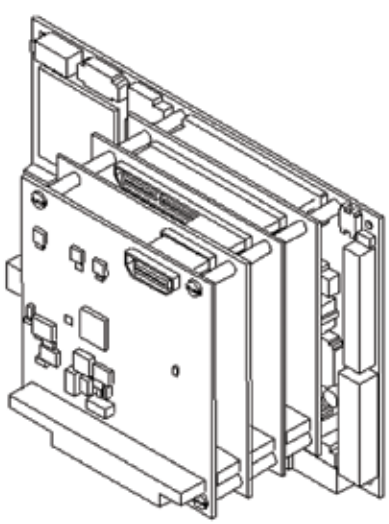

Figure 1 An EPIC Express Board with four stacking I/O modules.

Of the PCI Express hardware options on the market to date, none allow multiple boards with PCI Express to be stacked on top of each other. A stacking architecture is key for keeping costs low while maintaining a rugged, self-aligned group of boards. It also eliminates the need for a card cage or other external mounting frame.

Many questions arose about how to implement a PCI Express solution on EPIC. Should stackable, add-on legacy PC/104 I/O boards continue to be supported? If so, should both the PC/104 and PC/104-Plus expansion connectors remain? Is there a commercially available high-frequency connector that will allow PCI Express signals to be supported? If so, how many boards can be in a stack? Have the transmission line characteristics been modeled? Can a modular stack be defined that is position-independent yet requires no slot or address jumpers? Also, the enhanced capability of PCI Express is new and different from the legacy, parallel PCI bus architecture in that it

supports scalable link widths in 1-, 2-, 4-, 8-, 16- and 32-lane configurations. Which lane widths should be supported and how many channels of each?

The five embedded SBC manufacturers that defined and created EPIC (Micro/sys, Octagon Systems, VersaLogic, WinSystems and Ampro Computers) conducted extensive technical discussions to hammer out these issues in order to fully define EPIC Express. The result was a definition that offers both an evolutionary and revolutionary solution. It is evolutionary since it supports legacy I/O, yet revolutionary since it is the first stackable PCI Express solution. Its PCI Express architecture is a state-of-the-art serial interconnect technology that keeps pace with recent advances in processor and memory subsystems, providing a general-purpose interconnect of choice for a wide range of applications, including graphics, storage, networking, etc.

It is also important to note that the PC/104 connector was maintained because of the number of I/O modules available worldwide from hundreds of suppliers. For low-speed, simple I/O such as relays, digital I/O, low-speed communications and certain A/D functions, this is the easiest, lowest cost and most straightforward implementation for a system designer. They can either select an off-the-shelf unit or design their own.

An EPIC board with the addition of PCI Express becomes an EPIC Express board. None of the dimensions or I/O Zones of an EPIC board change. The only difference is the replacement of the parallel PCI bus (as implemented through the PC/104-Plus connector) with a serial PCI Express connector. Similarly, the I/O boards that stack on top of EPIC Express are PC/104-size with the PC/104-Plus connector removed and replaced with a PCI Express stacking connector. Some have dubbed these "PC/104 Express" modules (Figure 1).

EPIC Express Implementation

With PCI Express, the number of serial lanes to a device can be increased to address current and future bandwidth needs of I/O devices. The PCI Express Base Specification defines the configuration of serial links as x1, x2, x4, x8, x16 and x32. A PCI Express link can be scaled on a device-by-device basis to meet different I/O controllers' bandwidth and application objectives. It should be noted that the bandwidth of a x1 PCI Express link is nearly quadruple that of a 32-bit wide, 33 MHz PCI bus (as implemented with PC/104-Plus) with about one-fourth the number of pins. The EPIC Express technical committee considered both the application environment and practical implementation aspects of future bandwidth demands for small, industrial, stackable I/O modules and decided to support four x1 and two x4 links. However, one x16 link will be defined in the future to handle video applications.

There are two configurations that are currently defined and supported. The first is a replacement for existing, standard PCI devices using four x1 links (A-D). Similar to PC/104-Plus, it allows up to four boards to be stacked, yet requires only a single 28-pin connector. The connector is about 1/3 the size of a PC/104-Plus connector, which frees up valuable board real estate, quadruples the bandwidth of the link and eliminates the slot selection switch. This is called the Standard EPIC Express configuration.

The second configuration supports more bandwidth-intensive controllers and more PCI Express links. It supports the four x1 links on the first connector bank and the two x4 links (E &F), plus more clocks, power and ground. This adds two additional connector banks that take up approximately the same amount of area as a PC/104-Plus connector. This is called the Full EPIC Express configuration since it has more capacity and requires a larger connector with three 28-pin banks (Table 1).

The spacing between modules in an EPIC Express stack is maintained at 0.662-inches. This permits the same spacers and PC/104 cards to be used. The x4 EPIC Express expansion module must be closest to the root SBC. Stacked above that would be any x1 lane modules and finally

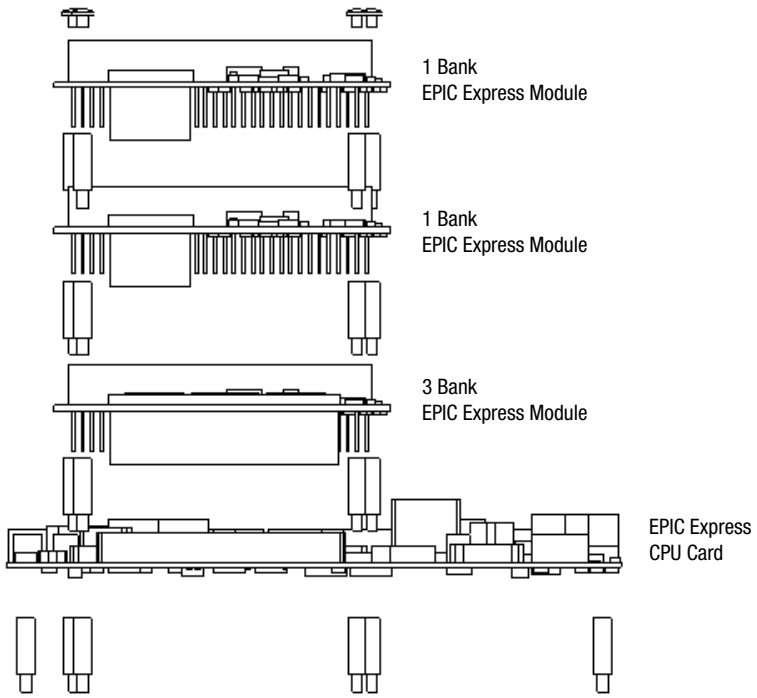

Figure 2 One-bank (standard configuration) EPIC Express Modules plug into 3-bank (Full configuration) EPIC Express Modules.

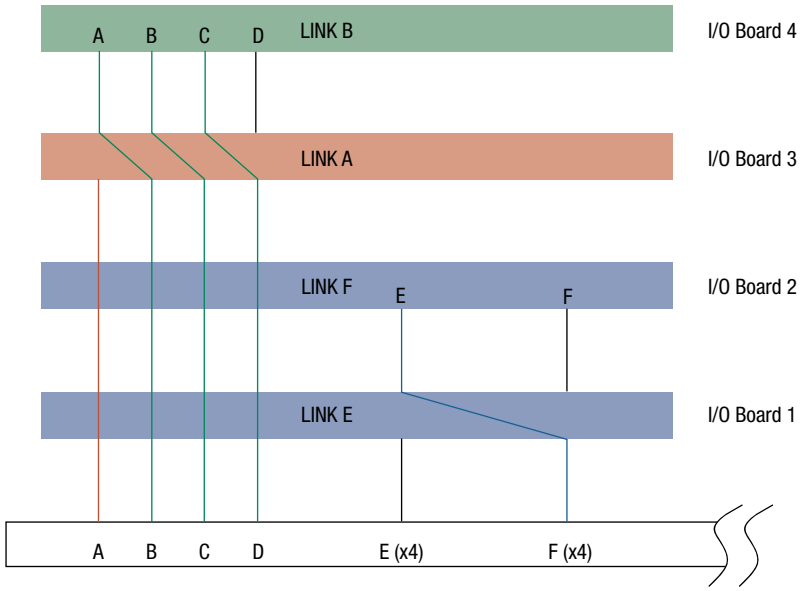

Figure 3 A four-board Full EPIC Express stack with two x4 links and two x1 links featuring automatic alignment for each.

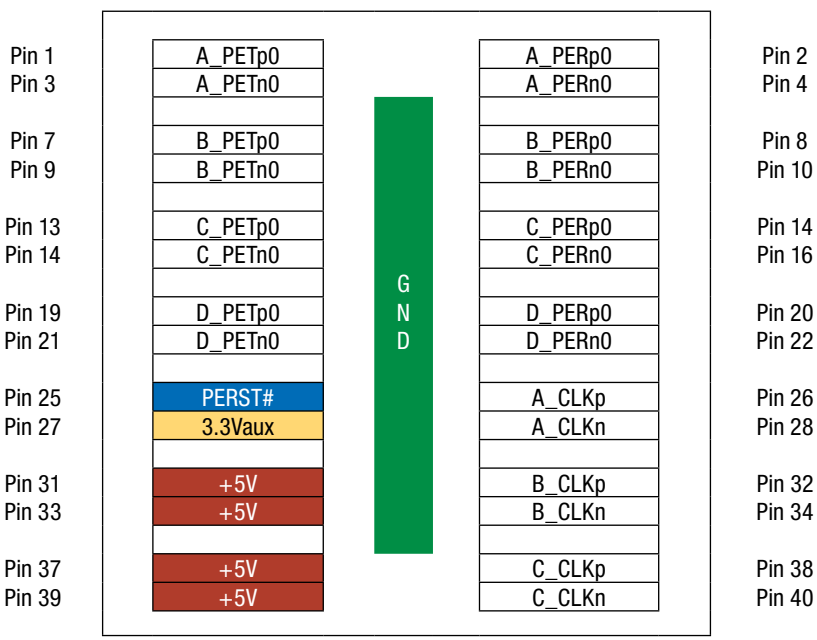

Figure 4 Pin Assignments for Standard EPIC Express Connector (Bank 1). An EPIC Express Standard connector duplicates the first third of the Full connector.

EPIC Express Configuration Options

	Standard	Full
Connector Size	1-bank	3-bank
Links	A, B, C, D	A, B, C, D, E, F
Lane Width	Four x1	Four x1, Two x4
Clocks	A, B, C	A, B, C, D, E, F
Power	+5V, PERST#, 3.3Vaux	+5V, PERST#, 3.3Vaux +12V, -12V

Table 1 EPIC Express configuration options.

PC/104-compatible modules. Even though the Standard and Full configurations require different EPIC Express connectors, a 1-bank, Standard EPIC Express Module can plug into a 3-bank Full EPIC Express Module but not vice versa. The reason is that a 3-bank connector is simply a 1-bank connector replicated two more times in a contiguous housing to support additional pins (Figure 2).

Automatic Link Alignment

One of the design goals of EPIC Express modules was that they not require any jumpers for address or slot alignment. The stacking design is physically similar to PC/104 and PC/104-Plus but the connectors employed are not through-hole style. Those were designed to implement a passive, parallel bus. By contrast, EPIC Express uses a pair of surface-mount connectors that allow one or more PCI Express controllers to be mounted on the I/O expansion module. This feature allows automatic link alignment, which eliminates the need for jumpers or a special stacking order. Boards not supporting PCI Express simply pass the signals up the stack from one connector to another.

An EPIC Express module with a single x1 PCI Express controller is always wired to Link A on the bottom side connector of its board. The top connector will have Link A wired from Link B on the bottom connector. The same methodology is maintained for Links B and C on the top connector since they are wired from Link C and Link D, respectively, on the bottom connector. Link D on the top connector is not connected.

An EPIC Express module with two x1 PCI Express controllers is always wired to Link A and Link B on the bottom side connector of its board. Therefore, Link A and Link B on the top connector will be wired from Link C and Link D, respectively, on the bottom connector. Link C and Link D on the top side connector will not be connected.

An EPIC Express module with a single x4 PCI Express controller is always wired to Link E on the bottom side connector of its board. The top connector will have Link E wired from Link F on the bottom connector. Link F on the top side connector will not be connected. Links A, B, C and D will be wired from the top

connector to the bottom connector since there is not a x1 PCI Express controller on the board (Figure 3).

EPIC Express Connector

PCI Express is a high-speed, low-voltage, differential serial interconnect that allows two devices to communicate with each other. It is designed to offer higher bandwidth with fewer pins. Each link has a full-duplex, transmit and receive pair of signals operating at a signaling rate of 2.5 Gbits/s. This speed of 2.5 GHz results in 2.5 Gbits/s for each direction that provides a 250 Mbyte/s communications channel in each direction (500 Mbyte/s total).

The connector is the enabling technology that allows PCI Express to be implemented in a stack. Samtec's Signal Integrity Division was instrumental in the definition and design process since they had done extensive work with their "Final Inch" differential connectors. They had worked with HyperTransport and were able to bring the same technology developed for their QTE/QSE connectors. They worked closely with the EPIC Express technical group and designed a special test rig and conducted testing to verify performance. The result is that a connector was specified that allows up to four boards to be stacked on top of an EPIC Express baseboard while maintaining proper electrical transmission line characteristics to support the data signaling rates.

Each connector bank is a 40-pin device with every third pin removed to support differential high-frequency signaling and a ground plane in the middle. The Standard configuration supports four x1 lanes and is a subset of the Full configuration. It is 1/3 the size of a Full connector, provides up to 16x the bandwidth of PC/104-Plus (parallel PCI bus) and has the same pin out for pins 1-40 (Figure 4). The Full configuration supports the four x1 and two x4 lanes on a single connector (Figure 5).

The assignment of the pin locations is important. They were selected to facilitate optimum routing of the PCI Express signal pairs as they are passed up the stack. Preliminary board layouts using the Intel-recommended PCI Express routing guidelines were conducted to

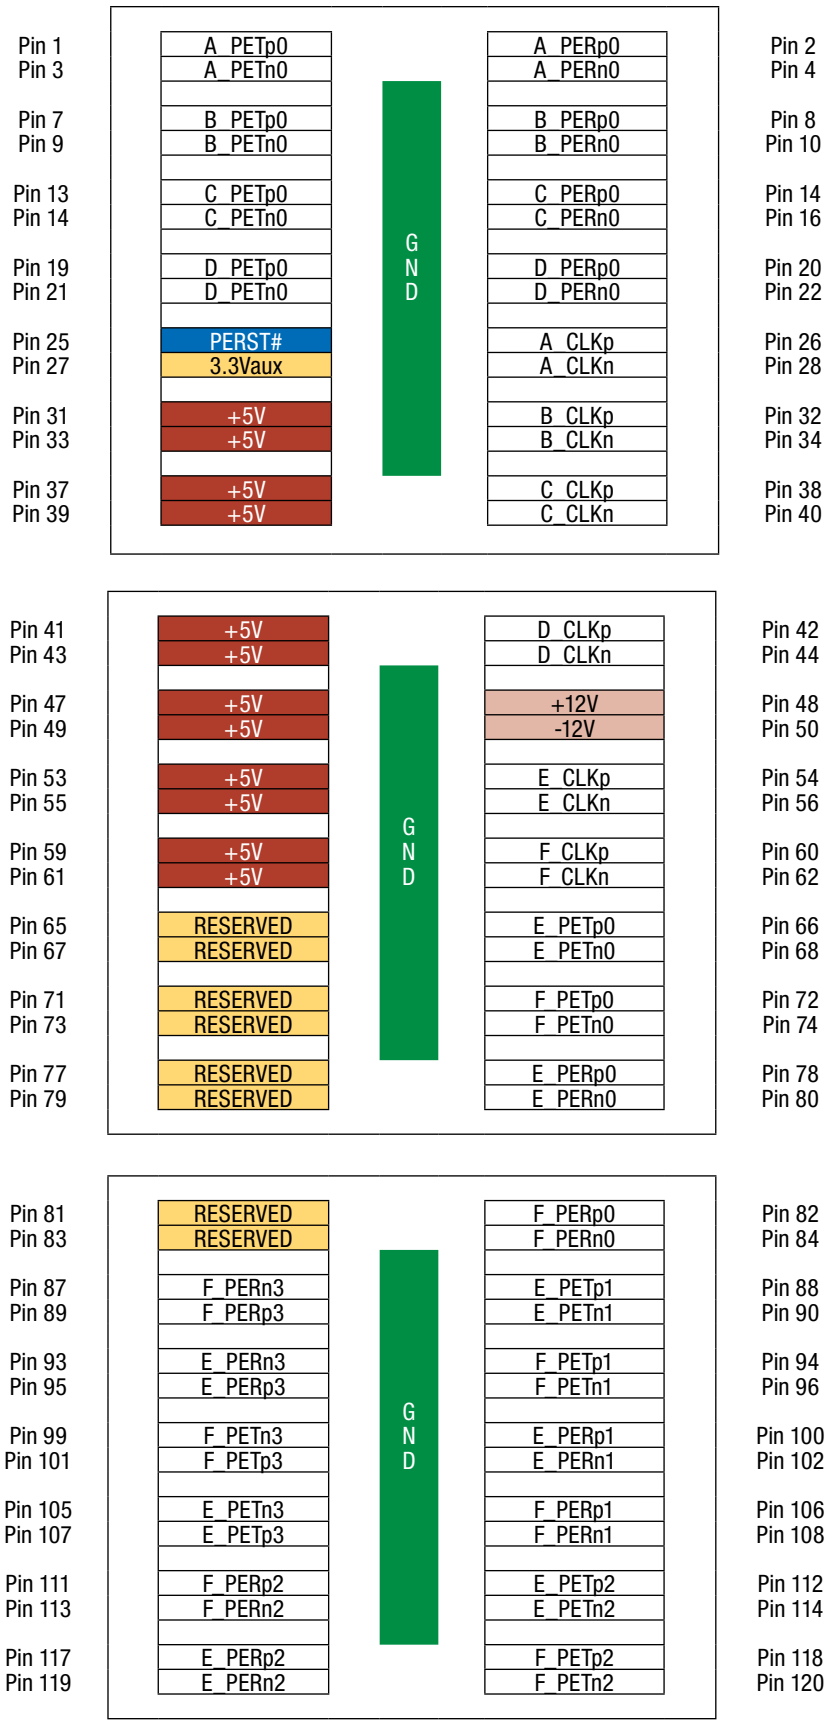

Figure 5 The EPIC Express Full connector supports four x1 and two x4 lanes.

SHARE YOUR VISION.

WE'LL SHARE OUR INNOVATION.

Share your vision with us. We'll customize our products to your requirements or partner with you to develop custom products all the way through deployment. Either way, you'll get leading-edge board level solutions for the most demanding embedded applications. Tell us what you need at **info@vadatech.com**.

- Comprehensive Hardware/Software solution for Intelligent Peripheral Management Interface (IPMI) version 2.0
- Symmetric Multi-Processing CPU for AMC/VME Modules
- A complete line of ATCA and AMC carriers
- High dynamic range A/D Converters

vadatech inc.
THE POWER OF VISION

www.vadatech.com 702.896.3337

ensure that the differential signaling environment could be maintained. This is necessary to minimize impedance mismatch, reflection and flight time, which degrade signal quality at these frequencies.

Generation 1 of PCI Express x1 links does not need a clock, but future generations double the serial data rate, which may require a clock, so pins are defined on the connector. Therefore, a clock is included for each link in the EPIC Express pin definition.

The transmit and receive pairs for the links, clocks for each link, power and reset are supported. However, PCI Express card Presence Detect and Wake are not. These two signals are for hot swap applications. Since EPIC Express expansion modules bolt together, it is not logical to implement hot swap. SMBus and JTAG are not currently supported. The +3V power rail is not supported since most all EPIC SBCs are +5V only. Therefore an EPIC Express expansion module will generate the requisite onboard voltage from the +5V rail.

On a Standard EPIC Express implementation, the first bank connector contains four x1 data links (A-D), three clocks (A-C), +5V, ground, auxiliary 3.3V and reset (PERST#). A Full EPIC Express implementation adds two x 4 links (E-F), three clocks (D-F), more +5V power pins, +12 volts, -12 volts plus four reserved pairs for possible functions to be determined in the future.

The designers of EPIC Express are delivering on their promise of upgrading to the latest technology. EPIC Express defines the addition of a PCI Express connector to both the EPIC SBC and PC/104-sized I/O expansion modules. The compact and incremental nature of EPIC SBCs using PC/104 expansion modules has proven beneficial in a wide range of embedded applications including test equipment, medical instruments, communications devices, transportation systems, military/COTS, data loggers, security, robotics, semiconductor manufacturing instruments and industrial control systems.

Unlike the other mid-sized boards, ample board space has been reserved to support the broad base of PC/104 I/O modules as well as PCI Express without sacrificing I/O capacity. Due to this architecture, compliant implementations can smoothly migrate from legacy ISA-based systems to the fast serial interfaces of the future. For more detailed information and a copy of the EPIC Express specification, visit www.epic-express.com.

WinSystems
Arlington, TX.
(817) 274-7553.
[www.winsystems.com].

AMCs Make Modular Multiprocessing a Reality

Modular computing has become a basic building block of embedded systems. From voice-over-IP packet processing to high-performance number crunchers, the ability to add processing units piecemeal provides upgradeability as well as scalability.

by Ben Klam
Extreme Engineering Solutions

Several standards have emerged to service the modular computing market. The Common Mezzanine Card (CMC) standard provides a small form-factor capable of supporting a CPU, system bridge, memory and flash. The PCI Mezzanine Card (PMC), a derivative of the CMC standard, added a PCI interface, allowing modules to communicate with one another and peripherals through a high-speed data bus.

The PMC provided a good platform for many types of computing systems. However, engineers were not satisfied with a single processor per PMC module. It was soon realized that there was just enough PCB real estate to support dual processors in a symmetric multiprocessing (SMP) architecture. In this environment, two processors share the same memory and I/O resources. For applications that mostly lived in the CPU caches, this architecture provided significant performance-versus-price improvements.

The PMC standard, however, did not account for the large power and cooling requirements of multiprocessor boards. PMCs were originally intended to host peripheral devices, such as SCSI and Ethernet controllers, not power-hungry CPUs. Additionally, features that were advantageous to high-performance computing, such as hot swap and board management, were not requirements when the PMC standard was created.

The Advanced Mezzanine Card (AMC) standard is the most recent addition to the suite of modular standards. AMCs increase the total amount of power allocated to each module, provide a larger selection of high-bandwidth interconnect, as well as support modular computing features such as hot swap and board management. Figure 1 shows a comparison of the two form-factors.

Figure 1 The PMC standard (left) provides a platform for modular computing. The recently ratified AMC standard (right) takes modular multiprocessing to a new level, while retaining a similar form-factor.

Multiprocessors

A multiprocessing system consists of multiple processing units tied to a single system controller bridge. This bridge usually consists of CPU interfaces, a memory controller and various I/O interfaces. As an example, consider the Marvell Discovery III PowerPC system controller. This embedded device consists of a 60x CPU bus interface, a DDR SDRAM controller, multiple gigabit Ethernet controllers, as well as two PCI bus interfaces (Figure 2).

The key advantage of a symmetric multiprocessing (SMP) system is the cost-effective sharing of resources such

EPIC Solutions for Real World Problems

Rugged, Reliable, and Ready-to-go

Based on the Embedded Platform for Industrial Computing (EPIC), the EPX-C3 combines the processor and I/O functions required for medical, transportation, instrumentation, communication, MIL/COTS, security, and other demanding applications.

- Fanless 733MHz C3 Processor
- Up to 2GB Flash memory
- 4x AGP CRT/LCD controller
- 10/100 Mbps Ethernet
- USB 2.0 support
- 4 COM channels
- 24 Digital I/O lines
- EIDE, FDC, and Kybd controllers
- PC/104 & PC/104-*Plus* expansion
- -40°C to +85°C operation
- Size: 4.5″ x 6.5″ (115mm x 165mm)
- Quick Start Developers kits for Windows® XP, CE, and Linux

Profit from our proven experience. We look forward to the opportunity to demonstrate how our success in the industrial market can work for you.

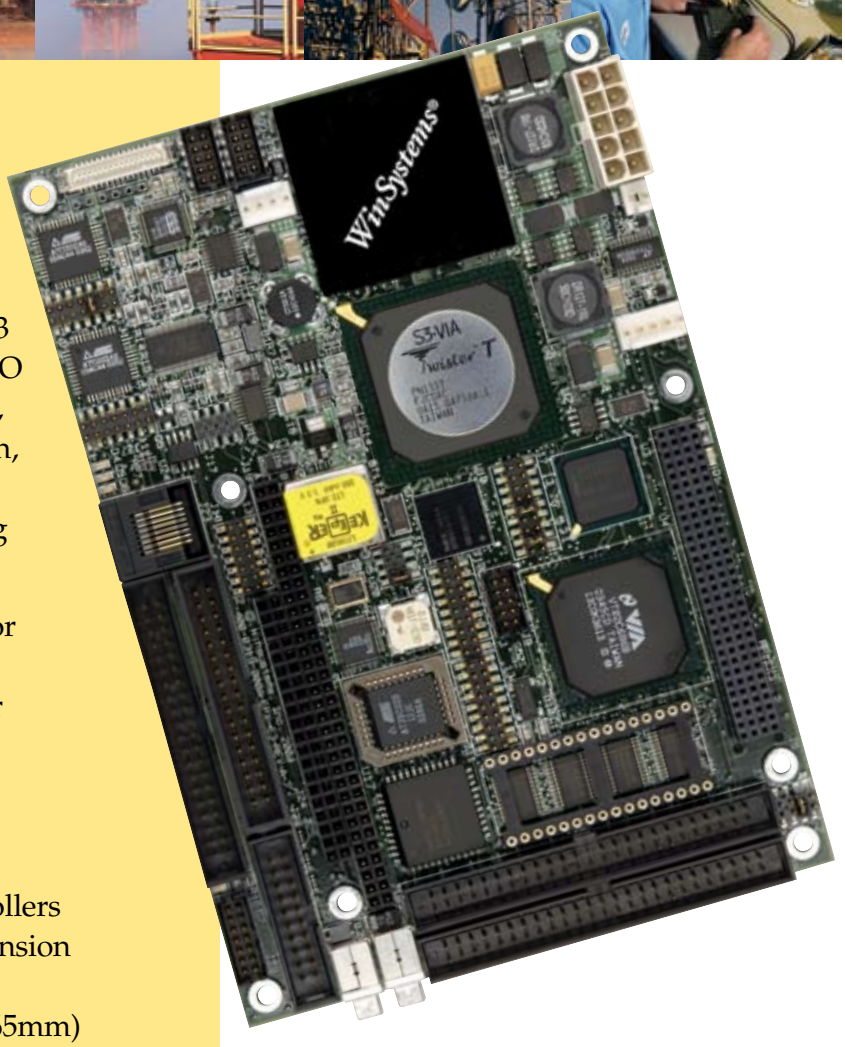

Call 817-274-7553 or
Visit www.winsystems.com

Ask about our 30-day product evaluation!

WinSystems®

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

as memory and interconnects. While it is true that each CPU in a dual-processor multiprocessing system only has half the potential memory bandwidth, there are a large number of computationally intensive applications that benefit from the reduced costs and PCB real estate of the multiprocessor approach. These applications typically are able to reside mostly within the caches of a processor. However, there are memory-intensive processor applications such as pattern searching that would be better suited to having a separate memory subsystem for each processor. One must perform a careful analysis of the application before considering an SMP platform over a uniprocessor approach.

SMP architectures also provide performance enhancements without the software headaches of clustering separate computer systems. In SMP, a single operating system is run on the processors. This places the majority of the load-balancing responsibility on the operating system and allows the programmer to follow a single processor model. The ease of sharing data structures is also a significant performance advantage of SMP. Shared data structures can be examined by any processor without the overhead associated with interconnect transactions or message-passing protocols.

PMCs

The PMC form-factor has several limitations when used in a multiprocessor system. These limitations include power dissipation and interconnect bandwidth, as well as lack of high-availability features.

PMCs are limited by the total amount of power they can draw from the 3.3V and 5.0V power rails provided. The power pins are limited to 1A per pin. This works out to 29.7W on the 3.3V rail and 30W on the 5V rail. While in theory this would yield around 60 watts of power, the CMC standard limits total power dissipation to 7.5 watts. The PrPMC (Processor PMC - VITA 32) specification extends this to a maximum of 25 watts.

Cooling a PMC can present additional challenges. The standard 10 mm and 15 mm module heights provide a limited amount of cross-sectional area for dissipating heat. For the 15 mm case, a cross-sectional area of 14 cm^2 is available for components, heat sinks and airflow. With a standard 200 linear feet per minute of airflow, it is difficult, if not impossible to dissipate more than 25 watts of power in a commercial 0° to 55°C operating environment.

PMCs utilize a PCI/PCI-X bus for interconnect with a carrier and other modules. PCI is a shared bus architecture, where each device presents an additional load on the common bus signals. The ability to meet signal integrity requirements limits the number of devices that can directly talk to one another at any given frequency. Transparent bridges can be used to create clusters of devices that operate at higher frequencies, with a tradeoff of higher transaction latency. When a point-to-point architecture can be used, PCI-X can be operated at 133 MHz. This yields a theoretical maximum data rate of around 1 Gbyte/s over a 64-bit bus.

PMCs are not hot-swap devices, making them unsuitable for the kind of high-availability systems that are beneficial to modular computing. They are directly mounted into the host carrier. However, in the case of a CompactPCI PMC carrier card, the entire assembly can be

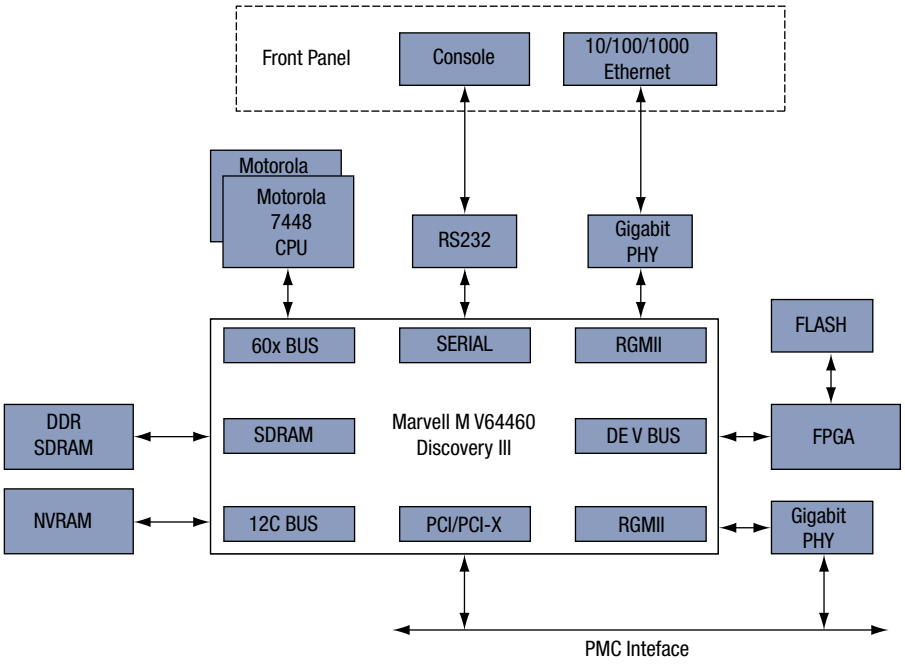

Figure 2 XPedite6200 by Extreme Engineering is an example of a multiprocessor PMC module. An operating system as well as file system can be stored within the on-card flash for stand-alone operation.

Figure 3 SMP Processor PMC modules require a novel approach to cooling. XPedite6200 utilizes a heat sink that covers both CPUs and the system controller.

hot-swapped. If the hot-swap carrier card is populated with more than one active module, then all modules much be taken out of service to replace the one faulty module. Clearly, this is not advantageous in a clustered modular computing environment. As an example, consider a PrPMC card consisting of dual Freescale 7448 1 GHz G4 PowerPC processors, a system controller, 1 Gbyte of DDR DRAM, 64 Mbytes of flash and two Gigabit Ethernet ports (Figure 3).

The first obstacle beyond fitting all the components in the PMC form-factor was satisfying the power requirements. It was a requirement that all power be dissipated from the 3.3V power rail, limiting the total power available to less than 25 watts. While a 25-watt maximum is allowed for a PrPMC, most carriers are designed to accept only PMCs, limiting the selection of carriers on which this PrPMC could be placed. Additionally, this power ceiling restricted the speed at which the CPUs could operate, forcing one customer to spend resources optimizing software to reach original goals.

The next challenge was cooling the CPUs, bridge, DDR SDRAM and Ethernet PHYs. Fortunately, a 15 mm stacking height form-factor was allowed, which yielded much more room than the typical 10 mm for a heat sink. This challenge required the creation of a novel heat sink, which covered both CPUs and system controller, allowing maximum surface area for cooling. However, the custom nature of the heat sink added unforeseen complexity to the final product.

Another PMC design issue was a requirement to have a watchdog timer perform a full hardware reset of the module, including the system controller. The PCI architecture includes a global reset asserted by the carrier card. This made it difficult to allow the module to act as a separate entity, capable of completely resetting itself without disrupting other members of the PCI bus.

Making the design fit into a PMC required several tradeoffs. These included specific PMC carrier card requirements, performance limitations due to power budgets and sacrifices in functionality. Had the design been based on the AMC standard, several of these tradeoffs could have been addressed.

AMC

As a more recent standard, the AMC architecture is able to provide several key advantages over PMC while maintaining a similar form-factor. Advantages include improved power capability, heat distribution, interconnect technology and support for high availability.

The maximum power consumption of an AMC is 60 watts. Instead of providing several separate power rails, AMCs consolidate all power onto a single +12V rail. Today's designs need to support ICs that operate at many different voltages. A typical design may need 2.5, 1.8. 1.5 and even 1.0V. Utilizing a single supply to derive all onboard voltages eliminates the need for board and system designers to balance power across multiple voltage rails, making power distribution much cleaner and easier to design.

The AMC standard also improves the ability to cool high-power components such as CPUs and bridges. The AMC specification provides several options for module height. A full height AMC module provides a cross-sectional area of 22 cm^2. This is a 57% increase in area as compared to a 15 mm PMC module. This additional area provides the necessary room for airflow and allows larger fins for greater heat sink surface area.

As processing performance increases, so does the I/O bandwidth. The AMC architecture supports three high-speed

Processor performance no longer needs to be sacrificed due to power and heat dissipation constraints. The 60-watt maximum AMC power easily allows the processors to run at maximum speed.

interconnect options: PCI Express, Gigabit Ethernet and RapidIO. Having three scalable options gives system designers much greater flexibility in meeting bandwidth requirements. As an example, PCI Express is a serial, lane-based interconnect that can support 1, 2, 4, 8, 16 or 32 lanes, each providing approximately 200 Mbytes/s of bandwidth.

AMCs also provide a important feature that PMCs lack: the ability to isolate and hot swap faulty modules. A malfunctioning PMC module requires removal from its PMC carrier card for servicing, and could even misbehave on the shared PCI bus, causing other modules or the entire system to malfunction. The point-to-point nature of the AMC interconnect standard allows a faulty module to be contained and removed for servicing without interrupting other system functions.

The AMC standard also includes an IPMI manager running on an auxiliary low-cost, low-performance microprocessor, and allows the AMC carrier to access information such as module power requirements, power supply health, temperature and other status. With this information, the AMC carrier can decide to shut down modules that are over-heating or about to fail. As an example, consider migration of the previous design to the AMC equivalent. The major changes include the migration from PCI/PCI-X to PCI Express and the addition of an IPMI controller.

Processor performance no longer needs to be sacrificed due to power and heat dissipation constraints. The 60-watt maximum AMC power easily allows the processors to run at maximum speed. The additional cross-sectional height also allows the use of standard heat sinks, saving the costs and complexity associated with exotic custom heat sinks.

The watchdog timer reset challenge that was identified in the development of the PrPMC module is no longer an issue. Since PCI Express is hot-swappable and point-to-point, the module can be locally reset, disabled or removed without disturbing the host and other modules. The AMC specification allows the full performance of the embedded SMP system to be realized. In addition, it adds features that improve the module's ability to satisfy the requirements of modular computing.

AMCs provide a feature-rich solution to today's multiprocessor needs. As the market continues to demand higher power, higher bandwidth and higher availability, the AMC standard provides room for SMP designs to grow.

While AMC modules are currently being deployed on AdvancedTCA carrier cards, PICMG, the developer of the AMC standard, is working toward a chassis and backplane standard. This system, called MicroTCA, will expand the possibilities of the AMC module and its ability to function as a building block to the next generation of cost-effective multiprocessing platforms. ◢

Extreme Engineering Solutions
Middleton, WI.
(608) 833-1155.
[www.xes-inc.com].

Next-Generation VME Boosts Defense and Aerospace Applications

With the addition of switched serial fabrics to VMEbus platforms, VXS can provide orders of magnitude more I/O bandwidth than is provided by legacy VMEbus systems.

by Robert Tufford
Embedded Communications Computing, Motorola

Since the turn of the century, it has become clear that the data bandwidth capabilities of parallel multidrop buses such as the VMEbus have been increasingly challenged in data-intensive applications. The result has been the launch of what is called the VME Renaissance—an era of VMEbus innovation and performance improvement while maintaining backward compatibility and protecting existing customer investments—to take VMEbus platforms to the next level of performance.

A series of stages were identified for the VME Renaissance. The goal of the first stage was to bolster the bandwidth capabilities of the multidrop VMEbus, and the goal of the second stage was to add switched serial fabrics to VMEbus platforms. The first stage came to fruition in August 2004 with the release of the Tundra TSi148 PCI-X-to-VMEbus bridge chip, which was co-developed with Motorola. The TSi148 implements the 2eSST VMEbus protocol, where "2eSST" stands for two-edge source synchronous transfer. The "source synchronous" part of the name indicates that the transmitter sends data along with a strobe. The "two-edge" portion of the name indicates that transfers are made on both the rising and falling edges of the strobe.

The 2eSST protocol enables products using the TSi148, working in conjunction with Texas Instruments' SN74VMEH22501 VMEbus transceivers, to provide up 320 Mbytes/s of data bandwidth in properly designed five-row backplanes. This is an 8X increase over the typical 40 Mbyte/s performance of the VME64 protocol. Motorola's MVME6100 and MVME3100, along with boards from other companies, use the TSi148 and provide a rich ecosystem of 2eSST-enabled products.

A use case example of how 2eSST technology is being considered to improve distributed application performance is the Navy Open Architecture initiative. This initiative has been focused on using Gigabit Ethernet, Data Distribution Service (DDS) and Common Object Request Broker Architecture (CORBA) middleware. 2eSST technology is now being evaluated for this initiative as a high-performance transport that complements Gigabit Ethernet. The next technology beyond 2eSST that could also be leveraged in this initiative is VXS.

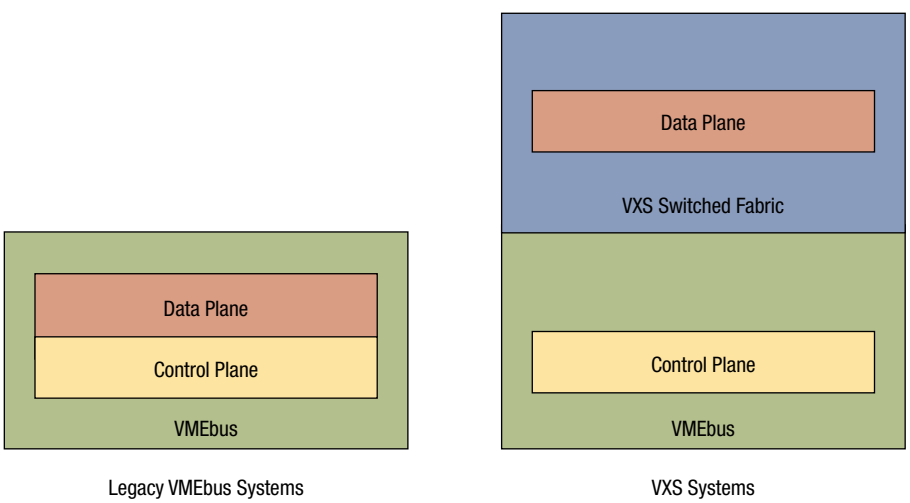

Figure 1 Data and control planes in legacy VMEbus systems and VXS systems.

VXS – The Second Stage of the Renaissance

The second stage formally began with the formation of the VXS (or VITA 41) working group in the VSO in March 2002. The goals for VXS as outlined at the inception of the working group were to:

- Provide a switched serial interconnect to VMEbus coincident with the VME parallel bus
- Utilize standard open technologies for the switched serial links
- Allow for multiple standard open technologies for the links, but not necessarily at the same time
- Preserve backward compatibility within the VMEbus ecosystem
- Increase the amount of power provided to each VMEbus card

Each of those goals has been achieved through the resultant specifications created by the working group:

VITA 41.0 - VXS Switched Serial Standard (Base Specification)

VITA 41.1 - VXS InfiniBand Protocol Layer Standard

VITA 41.2 – VXS RapidIO Protocol Standard

VITA 41.3 – VXS 1000 Mbit/s Baseband IEEE 802.3 Protocol Layer Standard

VITA 41.4 – VXS PCI Express Protocol Layer Standard

VITA 41.10 – Live Insertion System Requirements for VITA 41 Boards (Trial Use Standard)

VITA 41.11 – VXS Rear Transition Module (RTM) Standard

VITA 41.0-2 and 10 are currently moving toward ANSI standardization. The remaining standards will move to ANSI standardization in early 2006. These standards provide the framework for VXS implementations throughout the embedded computing industry, and a number of platforms have already been released based on them.

VXS – Data Plane and Control Plane

Now that the evolution and scope of the VXS standards have been characterized, how does VXS take VMEbus platforms to the next level of performance? Before tackling this question, it may be helpful to review the data plane and control plane concepts that are well known in the telecommunications industry. The data plane is defined as the portion of network traffic that is used to distribute data between nodes. The control plane is defined as the portion of network traffic used to set up, maintain and terminate data plane connections.

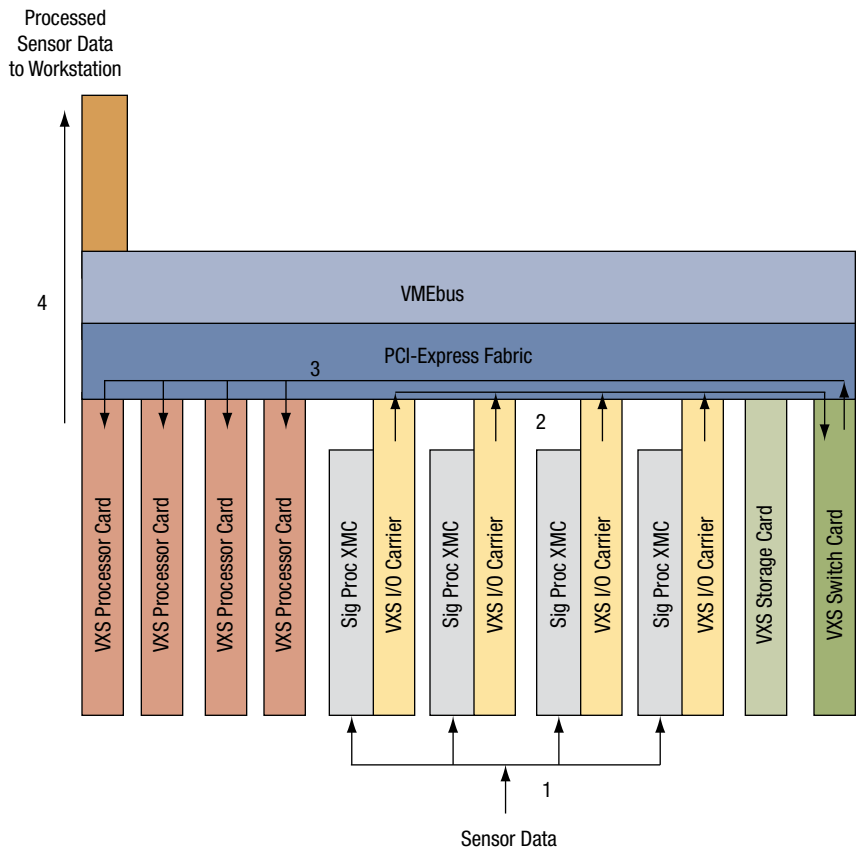

Figure 2 VXS application example: A signal processing system.

These concepts also apply to VME applications such as those found in the defense and aerospace markets. In legacy VMEbus systems, the VMEbus carries both control plane and data plane traffic (Figure 1). The VMEbus has always been well suited to carry control plane traffic, mainly because of its low latency. It also has been adequate to carry data plane traffic, provided that the application bandwidth requirements are less than 40 Mbytes/s (for VME64). However, one of the drawbacks of the VMEbus has been the fact that, since it is a multidrop bus, only one connection between nodes can occur at a time, so the aggregate bandwidth of the VMEbus platform is still only 40 Mbytes/s. The closest thing to multiple connections occurring on the VMEbus is broadcast mode, in which one node sends a message to multiple other nodes simultaneously.

Adding switched serial fabrics to VMEbus platforms, through VXS, overcomes the inadequacies of the VMEbus for data plane traffic in two ways. First, switched fabrics tend to have much higher data bandwidth capabilities than the VMEbus. For example, a 4X PCI Express link has a raw data bandwidth of approximately 2 Gbytes/s, or 1 Gbyte/s [(2.5 Gbit/s x 4)/10] in each direction, compared to the 40 Mbytes/s for the VMEbus. Second, switched fabrics use point-to-point connections, which allow multiple connections to occur simultaneously. As an example, the aggregate bandwidth for

a VXS 10-slot system, with 4X PCI Express links between five pairs of cards each exercising a 2 Gbyte/s simultaneous connection per pair, could be as high as 10 Gbytes/s, whereas the aggregate VMEbus bandwidth is still only 40 Mbytes/s.

Thus, VXS preserves the multidrop VMEbus for use in legacy applications, where it functions as both the data and control plane, while also allowing the data and control planes to be split between VMEbus and switched fabrics respectively to address data-intensive applications (Figure 1).

VXS Application Example

Consider as an example a signal processing system implemented using VXS (Figure 2). It consists of a 10-slot VXS chassis and backplane with the VMEbus routed to each slot, as well as a single-star PCI Express fabric routed from the VXS switch card slot to every slot. The chassis is populated with four VXS processor cards, four VXS I/O carriers with a signal processing XMC on each carrier, a VXS storage card, a VXS switch card and an RTM opposite the first VXS processor card. The VXS storage card is used for booting the system and storing configuration data.

The data flow is shown by the red arrows. Sensor data enters the system through the front panel ports of the four signal processor XMC cards (the arrows marked 1) and is preprocessed by the DSP engines on these cards. The preprocessed data is then routed through the VXS I/O carrier cards, across the PCI Express fabric, onto the VXS switch card (arrows marked 2). Note that each of the VXS I/O carrier cards can transfer data at up to 1 Gbyte/s simultaneously. This equates to an aggregate I/O bandwidth of 4 Gbytes/s—two orders of magnitude larger than is achievable with VME64! The data is then parsed by the VXS switch card to the appropriate VXS processor cards for post-processing (arrows marked 3). The first VXS processor card then sends out the processed sensor data to an external workstation for viewing (arrow marked 4).

The instantiation of 2eSST and VXS standards-based COTS technologies in VMEbus platforms brings rich benefits to defense and aerospace applications. 2eSST technology can boost the bandwidth of the legacy VMEbus by up to 8X. These performance increments are especially beneficial for defense and aerospace applications that require systems that are able to capture, process and deliver real-time data for display and control of actuating systems. VXS also allows for flexibility in configuring systems with the right balance of processing power and I/O ports, because the I/O bandwidth provided by the switched fabric interconnect now allows the system to behave as if it were one large virtual card. ◢

Motorola
Tempe, AZ.
(800) 759-1107.
[www.motorola.com/computing].

Built to go in harm's way.

NEW! TPPC64™ (IBM PowerPC 970FX-based SBC)

6U VME64 engine with up to two IBM® 1.8 GHz PowerPC® 970FX CPUs
Up to 4GB of DDR400 SDRAM
VME interface - Universe II VME64x interface
Two 10/100/1000 Ethernet ports
Six USB ports
AC97 audio, two serial ports
Support for Linux™ OS
Ruggedized for up to 30G shock

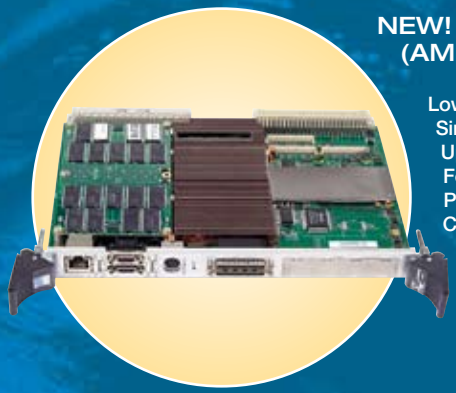

NEW! TA64™ (AMD Turion 64 Mobil-based SBC)

Low Power, AMD® Turion™ 64 Mobile CPU
Single/2-slot VME configurations
Up to 4GB ECC DDR memory
Four PMC slots, two USB ports
PS/2 keyboard and mouse port
Compatible with USPIIe-USB™ SBCs
64-bit Solaris 10, Windows® and Linux support
Ruggedized for up to 30G shock

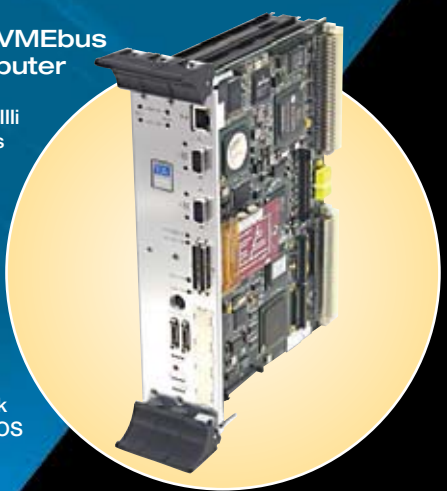

USP IIIi™ VMEbus Computer

Two UltraSPARC® IIIi 1.2 GHz CPUs
4GB 266DDR SDRAM per CPU, 8GB total
Optional TGA3D+™ graphics
Gigabit Ethernet port
One 64-bit/66-MHz PMC slot
Up to three PMC expansion slots
Dual FC-AL ports
Additional VME-backplane access
Up to 30G shock
Solaris™ OS

Our ultra-rugged, mission-critical SBCs now support more technologies than ever.

Extreme environments that need mission-critical reliability need Themis' single-board computers.

Our high performance, high availability computers have proven themselves in the harshest, most demanding applications and operating environments. Where other systems fail, our systems survive – in extreme temperatures and up to 30G shock.

And now Themis' family includes AMD and IBM PowerPC processors in addition to our well-known UltraSPARC VME and Compact PCI single board computers and graphic accelerators. So we can support applications in Solaris, Windows, Linux, UNIX®, and more.

All are designed for maximum flexibility and low Total Cost of Ownership. Because the budgeting process can be a pretty hostile environment, too.

So when your mission takes you in harm's way, you can rely on Themis.

www.themis.com (510) 252-0870

THEMIS

Transformational.

VITA 42 XMC Gains Momentum, Increases Flexibility

As I/O requirements continue their relentless upward progression in both performance and complexity, system integrators need a new I/O module architecture. VITA 42 reproduces the open standard ecosystem of the PMC module while providing higher performance and the potential for direct connection to switched fabric interconnects.

by Andrew Reddig
TEK Microsystems

High-performance embedded systems have always tried to balance the need for raw processing power with the ability to move data—at high speed—in and out of the system. The VITA 42 XMC standard is directly targeted at the requirements of high-performance systems in a rugged, deployed environment. Based on the successful PMC standard, XMC adds up to two high-speed serial connectors (designated J15 and J16) that can be used for switched fabric solutions as well as point-to-point connections between I/O modules and carrier cards. XMC leverages the availability of merchant silicon solutions for fabric endpoints and switches while maintaining the form-factor and mechanical characteristics to support operations in severe environments.

The XMC standard defines a layered architecture, with the base standard—VITA 42.0 —defining the mechanical, power and utility interfaces while several protocol standards map specific switched fabric interfaces into the XMC architecture. The current set of protocol standards are shown in Table 1, along with

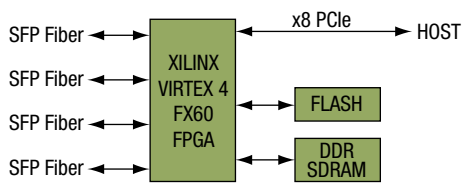

Figure 1 FPGA-based XMC module (4-channel Serial FPDP or Gigabit Ethernet interface using FPGA IP cores).

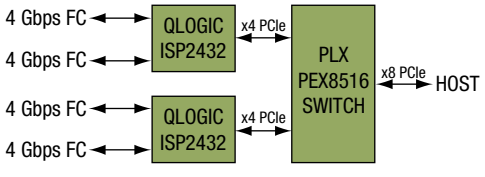

Figure 2 ASIC-based XMC module (4-channel Fibre Channel interface using QLogic ISP2432 controllers and PLX 8516 PCI Express switch).

the throughput for interfaces with one or both XMC connectors.

The XMC architecture has many features in common with another standard, the Advanced Mezzanine Card, or AMC, being developed by PICMG as a part of the AdvancedTCA ecosystem. Both XMC and AMC support direct connection to switched fabrics and leverage the growing number of endpoint solutions available from silicon providers. While the architectures are similar, the mechanical features are quite different due to the markets addressed by the two standards. AMC modules are designed to meet the needs of the high-availability telecom market, which requires that modules be hot-swappable without powering down the system or removing the carrier card from the chassis. This design goal resulted in selection of a card-edge-type connector that meets the requirements of the telecom environment but is unsuitable for the shock, vibration and humidity requirements of a deployed military application. XMC modules, on the other hand, have no hot-swap requirement and use a traditional stacking connector between the mezzanine and carrier that has been characterized for rugged environments.

Performance Beyond Limits

... ready for Extreme Embedded Computing?

ETXexpress

Reliable COM.0 Solutions

INSERT-READY

ACCELERATED DESIGN

SCALABILITY

LONGEVITY

COST-EFFECTIVE

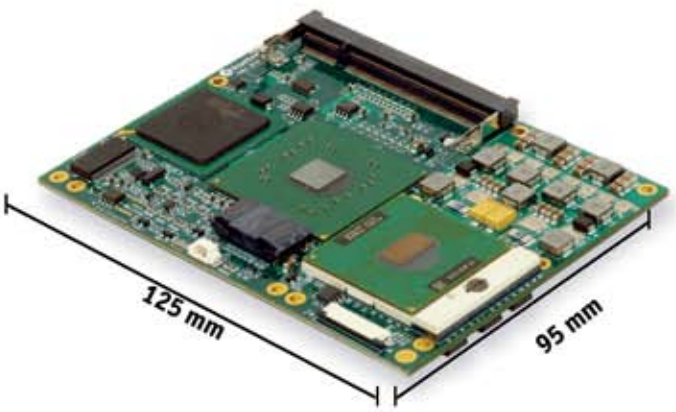

ETXexpress-PM

- Highest performance state of the art embedded module
- Intel® Pentium® M processor and advanced Intel® chipset

ETXexpress products are next generation embedded modules based on the PICMG COM Express standard. ETXexpress provides the hightest performance and I/O bandwidth available in COMs.

- **PCI Express - the elemental data path**
- **Gigabit Ethernet - for high connectivity**
- **USB 2.0 - for fast periphery**
- **Serial ATA - for fast drives**
- **ACPI - for optimized power management**

COM Express

ETX express™ www.etx-express.com

Intel® Communications Alliance Premier Member

PICMG

Extended Temperature Range ETR

Get ready. Get ETXexpress

Visit www.kontron.com/ETXexpress

Computer-On-Modules | Blades & Mezzanines | CPU Boards | Systems | Mobile Rugged | Custom Solutions

kontron

VME PRODUCT SHOWCASE

Featuring the latest in VME Product technology

T2-6U-VME – Octal ADSP-TS201 TigerSHARC 6U VME board

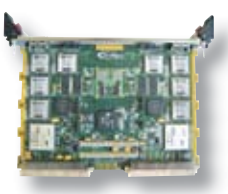

- Eight ADSP-TS201S TigerSHARC® DSPs providing 28.8 GFLOPs and 115 BOPS of processing
- Two Xilinx Virtex-II Pro FPGAs – interfacing, coprocessing
- ATLANTiS architecture – 8 GB/sec of external I/O throughput via Virtex-II Pro
- One PMC site w/PMC+ extension
- Tundra Tsi148™ PCI-X-VME bridge w/2eSST support
- 64-bit 66 MHz PCI interface via BittWare's SharcFIN PCI-DSP bridge

BittWare, Inc.
Phone: (603) 226-0404 **E-mail:** info@bittware.com
Fax: (603) 226-6667 **Web:** www.bittware.com

VME Backplane Product Family

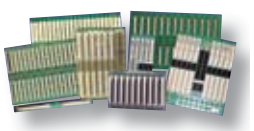

- Largest selection of VME-based backplanes in industry VME, VME64x, VME320, VXI, VXS (V41), Gig E (V31.1), VPX (V46)
- Leader in switched fabric implementations
- Optimized via signal integrity studies
- Specialize in custom designs
- Controlled impedance stripline design
- Hybrid versions available

Elma Bustronic
Phone: (510) 490-7388 **E-mail:** info@bustronic.com
Fax: (510) 490-1853 **Web:** www.bustronic.com

VME System Platforms

- VME, VME64X and VXS chassis, backplanes and accessories
- Rackmount, desktop, and portable versions
- Modular platform allows limitless configurations
- Advanced EMC shielding and system monitoring optional
- Designs optimized using backplane signal integrity studies and thermal simulation
- Customization available

Elma Electronic
Phone: (510) 656-3400 **E-mail:** sales@elma.com
Fax: (510) 656-3783 **Web:** www.elma.com

Tundra Tsi148™ PCI/X to VME Bridge

- Industry's best sustained transfer rate of 305 MB/s in 2eSST - higher system bandwidth
- Multithreading - capable of a number of simultaneous transactions while optimizing bus utilization
- Full featured master, slave and system controller - can be used in any VME application
- PCI-X local bus supports two loads at 133 MHz - reducing component count
- Small device footprint - saves board space
- Proven VME backwards compatibility - preserves legacy investment

Tundra Semiconductor Corporation
Phone: (613) 592-0714 **E-mail:** sales@tundra.com
Fax: (613) 592-1320 **Web:** www.tundra.com

Standard	Description	Throughput (J15 only)	Throughput (J15 / J16)
42.1	Parallel RapidIO	8-bit 500 MHz (1.0 Gbyte/s)	16-bit 500 MHz (2.0 Gbyte/s)
42.2	Serial RapidIO	8 x 3.125 Gbit/s (2.5 Gbits/s)	16 x 3.125 Gbps (5.0 Gbyte/s)
42.3	PCI Express	8 x 2.5 Gbps (2.0 Gbyte/s)	16 x 2.5 Gbps (4.0 Gbyte/s)
42.4	HyperTransport	8-lane 800 MHz (1.6 Gbyte/s)	16-lane 800 MHz (3.2 Gbyte/s)
42.5	Aurora	10 x 3.125 Gbps (3.1 Gbyte/s)	20 x 3.125 Gbps (6.2 Gbyte/s)

Table 1 XMC Protocol Standards

The evolution of the XMC ecosystem has gone through several stages based on the availability of switched fabric endpoints. The initial wave of activity was based on Parallel RapidIO, or VITA 42.1, driven by the availability of Parallel RapidIO processors from Freescale and switches from Mercury and Tundra. The next wave of activity leveraged FPGA technology to implement PCI Express (VITA 42.3) and Serial RapidIO (VITA 42.2) endpoints, both of which can be implemented using a common hardware platform and FPGA IP cores. Unlike the Parallel RapidIO and HyperTransport (VITA 42.4) standards, serial switched fabrics such as PCI Express and Serial RapidIO share a common physical interface, allowing the development of fabric-agnostic hardware products using FPGAs at one or both ends of the link. The ability to "hedge" the fabric wars has made it easier to move forward with XMC products without having to choose a specific fabric, as long as the application supports the power and cost tradeoffs of an FPGA-based design. An example block diagram of an FPGA-based XMC module is shown in Figure 1.

As the base of FPGA-based products reached critical mass, application developers realized that the FPGA could often be better utilized as a signal or protocol processing resource rather than using a large fraction of the available gates for a fabric endpoint. This resulted in applications being developed based on the Xilinx Aurora protocol as a lightweight, point-to-point link between XMC modules and carriers, or between different carriers on a VXS backplane. Although Aurora shares the same physical interface as higher level switched fabrics such as PCI Express or Serial RapidIO, the protocol is much more

XMC in a Data Recording Application

One type of application that benefits from the bandwidth and scalability offered by XMC modules is high-performance data recording and playback. One specific application currently being developed at Tekmicro on a PMC card, but which will be deployed on the new JazzFiber XMC card, targets recording and playback of multiple Serial FPDP and Gigabit Ethernet I/O streams using a Storage Area Network (SAN) array using 4 Gbit/s Fibre Channel.

The system level requirement is to record a mix of Gigabit Ethernet and Serial FPDP input streams and to reproduce the recorded data with precisely aligned timing during playback. Each sensor input requires individual packets to be time-stamped using a system-wide reference with a channel-to-channel accuracy of 100 ns. The packets are then buffered and stored to multiple channels of disk storage using Fibre Channel interfaces to either dedicated RAID controllers or SAN volumes. The system is required to automatically self-configure based on installed hardware to support different operating modes with the same basic hardware and software architecture. A Web-based interface allows the system to be controlled and monitored from any type of workstation without client-side software.

The Gigabit Ethernet and Serial FPDP interfaces are both implemented with Tekmicro's four-channel JazzFiber XMC module, shown in Figure 2 of the main article. By using an FPGA-based module for the sensor input and output streams, the system can perform application-specific channel processing, time-stamping and windowing of the data in real time. The module can adapt to any mix of Gigabit Ethernet and Serial FPDP channels through FPGA IP reconfiguration, allowing the same hardware to be adapted to multiple mission scenarios. In Serial FPDP mode, the XMC module supports just under 1 Gbyte/s sustained throughput across four channels using a x8 PCI Express interface at the XMC connector.

The storage interface is implemented using a quad Fibre Channel XMC module, shown in Figure 3 of the main article. This interface uses standard Fibre Channel protocols, so the optimum solution is an ASIC-based module using off-the-shelf silicon rather than the higher power and cost of an FPGA-based solution. The Fibre Channel XMC module supports four channels of 4.25 Gbits/s for a total throughput of up to 1.7 Gbytes/s. This module also uses a x8 PCI Express interface, which provides 2.0 Gbytes/s of bandwidth between the XMC and the carrier card.

The use of XMC-enabled carrier cards with onboard PCI Express switching enables a scalable implementation that can utilize between 1 and 16 carrier cards in a single chassis. With a single carrier card, the architecture supports 1 Gbyte/s sustained throughput between sensor I/O and storage, while a 16-card system supports up to 20 Gbyte/s aggregate throughput with up to 48 RAID interfaces. The use of XMC's increased bandwidth along with a direct fabric interconnect allows this single chassis system to scale to 3x higher throughput than previously possible with PMC technology.

XMC ultimately offers the systems integrator a wide range of compatible, interoperable I/O solutions with the performance and flexibility needed by tomorrow's high-performance embedded systems.

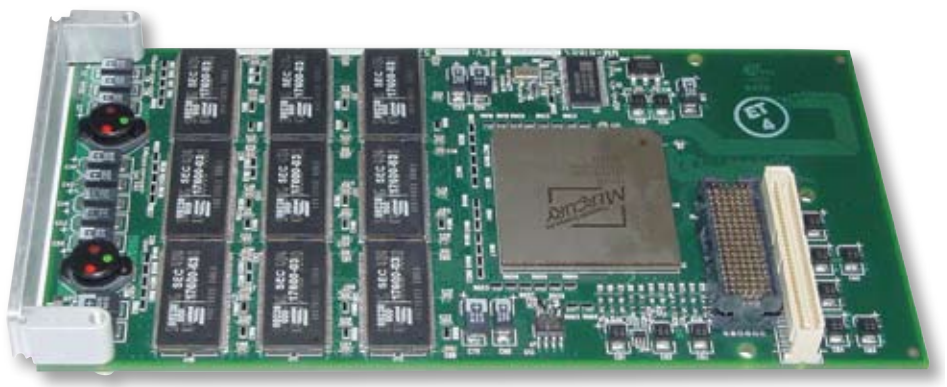

Figure 3 Mercury's System Memory XMC connects 2-4 Gbyte SDRAM directly to a RapidIO fabric.

Company	Protocols Supported	Description
Tekmicro	PCI Express, Serial RapidIO, Aurora	Quad Serial FPDP Quad FPGA-based fiber optic I/O Quad Fibre Channel
AdvancedIO	Parallel RapidIO	10 Gigabit Ethernet
GDA	Parallel RapidIO	8540-based processor module
Mercury	Parallel RapidIO	Dual Serial FPDP Bulk Memory Serial RapidIO chassis-to-chassis bridge
Pentek	PCI Express, Serial RapidIO, Aurora	A/D, FPGA
Spectrum Signal	Parallel RapidIO	A/D, D/A, FPGA
VMETRO	PCI Express	Bus Analyzer

Table 2 XMC Modules

lightweight, supporting a multi-lane link between two endpoints without routing or other system-level features. Unfortunately, Aurora by itself does not define enough details to provide an interoperable standard across different vendor implementations, which has resulted in multiple vendor-specific implementations with limited compatibility.

To address this, the VSO recently kicked off a new standard task group, VITA 55, to develop an Aurora-based lightweight protocol and address the interoperability issues. Two implementations of VITA 55 are being standardized: VITA 41.5 for VXS systems to support carrier-to-carrier links and VITA 42.5 for XMC to support mezzanine-to-carrier links. Fabric-agnostic implementations of both VITA 41 and 42 products are expected to support VITA 55 as one option along with existing PCI Express and Serial RapidIO implementations, offering a range of choices to system implementers to tradeoff fabric complexity with FPGA resource utilization.

Although FPGA-based solutions offer a high degree of flexibility, there is a penalty in terms of both cost and power. The next wave of XMC activity will leverage the availability of switched fabric endpoint silicon to implement standard I/O functions in XMC modules. While these solutions will not have the flexibility or fabric-agnosticism of an FPGA-based solution, they will provide high throughput and performance at much lower power and cost.

In particular, PCI Express-based modules for multi-channel Fibre Channel, Serial ATA, high-performance graphics, Gigabit Ethernet and 10 Gigabit Ethernet are under development by several vendors. A four-channel Fibre Channel XMC module implemented using PCI Express silicon is diagrammed in Figure 2.

Several companies have announced products based on the XMC standard (Figure 3), and more are expected later in 2005 and in 2006. Table 2 lists some of the XMC modules available from different vendors, and Table 3 lists some of the XMC carrier card options for various form-factors.

The XMC ecosystem is growing rapidly, with both highly flexible FPGA-based

Company	Protocols Supported	Description
Tekmicro	PCI Express Serial RapidIO Aurora	6U VITA 41 (VXS)
Curtiss-Wright	PCI Express Serial RapidIO	6U VITA 41 (VXS) 6U VITA 46 (VPX)
Interface Concept	Parallel RapidIO	6U VME64x
Mercury	Parallel RapidIO	PowerStream 7000 (proprietary) PICMG 3.0 (AdvancedTCA)
Micro Memory	Serial RapidIO	6U VITA 41 (VXS) 6U VITA 46 (VPX)
Pentek	PCI Express Serial RapidIO Aurora	6U VITA 41 (VXS)
Spectrum Signal	Parallel RapidIO	PCI-X

Table 3 XMC Carrier Cards

solutions and a range of lower power ASIC-based solutions available to meet the needs of systems developers. As the switched fabric silicon ecosystem expands, the range of XMC options will increase both for existing fabrics such as PCI Express and Serial RapidIO as well as future choices such as Advanced Switching and HyperTransport. FPGA-based solutions will continue to offer higher performance and flexibility for applications that need tailored solutions, with next-generation FPGAs from Xilinx, Altera and others offering better performance at lower power and cost.

XMC ultimately offers the systems integrator a wide range of compatible, interoperable I/O solutions with the performance and flexibility needed by tomorrow's high-performance embedded systems.

TEK Microsystems
Chelmsford MA.
(978) 244-9200.
[www.tekmicro.com].

Deploying VITA 46 in Real-World Applications

The embedded market will soon witness the debut of the first boards and systems based on the emerging VITA 46 architecture, an industry milestone that promises to deliver the next generation of embedded computing board performance.

by John Wemekamp, Stephane Joanisse and Jing Kwok,
Curtiss-Wright Controls Embedded Computing

Designed to address the performance and environmental limits now confronting VME64x, VITA 46 provides an open standard for hardware form-factor and electrical connectivity uniquely suited for high-bandwidth distributed processor systems. It provides several significant advances over traditional VME including its high-speed connector interface and, through the complementary VITA 48 standard, support for 2-Level Maintenance for in-the-field repair and replacement. Three particular application case examples highlight the most likely target applications for early use of VITA 46: these include radar systems, mission computer systems and small form-factor graphics display systems. Each of these applications highlights specific technological advantages of VITA 46 over the existing VME64x or CompactPCI systems that it will replace or enhance in heterogeneous systems. Table 1 shows a comparison between VME64x and VITA 46.

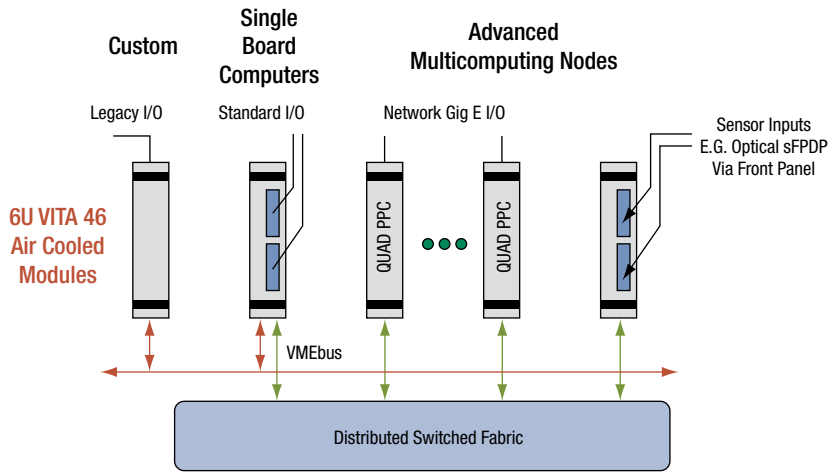

Figure 1 Radar processing—VITA 46 enables theuse of next-generation switched fabrics such as Advanced Switching Interconnect, Serial RapidIO, etc., in either centralized or distributed architectures.

Radar Processor Systems

Radar systems comprise a classic "hard" problem for embedded multicomputing: they typically involve multiple channels of high-speed streaming input data. Significant data errors can occur if the streaming data is interrupted to allow processing to catch up. In radar systems, performance must be as close to real time as possible. As latency in the system increases, so does the risk of data loss. Because human operators use the processed data from these systems for targeting or navigation, the resulting data can be life-critical, making it all the more important to minimize latency as much as possible.

Today's VME64x-based radar processing systems typically receive their

The Power of Choice...

Command the game with your next I/O move.

Modularity. Scalability. Reliability. Cost-effectiveness. These represent the solid foundations that SBE delivers to OEMs for building innovative end solutions. Partnering with SBE for networking and communications I/O solutions allows you to take advantage of proven technology and field-tested products designed to optimize performance for your unique application needs.

SBE offers a full spectrum of interface cards, ranging from T1 and T3 to Gigabit Ethernet and IPsec/SSL acceleration. These boards are available in multiple form factors, including PCI, PMC, and PTMC. Customers have the choice of buying these boards individually or bundling any of the PMC/PTMC modules with our intelligent core processing platforms to create a flexible, cost-efficient blade solution ideal for serving demanding telecom applications. Full Linux support is available on every board.

- Channelized T3
- 24-port T1/E1
- LAN/Ethernet
- Storage
- IPsec/SSL Encryption
- Blade platforms
- I/O and beyond...

sbe®

Linux *On Demand*

flexibility on demand | 925-355-2000 | info@sbei.com | www.sbei.com


YOU LEFT NOTHING OUT
OF YOUR NEW PRODUCT.
Except 18 months of development time.


datacom | telecom | wireless devices | automotive | medical | avionics
When speed to market is the driver, leading product developers turn to one company for help: Enea. Why? As the world leader in embedded communications, our high-availability OSE™ RTOS significantly shortens the design cycle time while enhancing system reliability. So you'll not only save money and get to market faster, you'll get there with a better product. To learn more about how our professional grade software and solutions can shorten your time to success, call 866-844-RTOS or go to enea.com.
ENEA
Embedded for leaders.

Headless, But Not Mindless.

Add Micro/sys brainpower to your embedded application.
There's a lot riding on that new OEM design: sophisticated application code, tricky I/O interfaces, and system integration issues. You don't need a headache with CPU design issues, too. You just want a reliable CPU that will plug into your application, interface with your system I/O, and run. Headless might be the right way to go, but the price has got to be right.

We have just the processing brainpower you need!
For example, the new Micro/sys SBC1586 is a headless PC/104 size Pentium® workhorse. It includes 10/100BASE-T Ethernet, USB, four COM ports and a removable CompactFlash® for data logging or OS. This embedded PC supports Linux, DOS, Windows® NT or CE, and VxWorks.® Starting under $500.

For a less scary ride on your development adventure, choose from a wide selection of X86 through Pentium CPUs, headless or fully configured with A/D, D/A, and other PC I/O. Call today and let the minds of our engineering and customer service teams help you rein-in those unbridled development costs by offering a compact and cost-effective solution.

MICRO/SYS

3730 Park Place, Montrose, CA 91020
Voice (818) 244-4600 Fax (818) 244-4246
info@embeddedsys.com

www.embeddedsys.com

For more fully-featured Micro/sys embedded PCs in various sizes, options and interfaces, take a look at the boards below and check our Web site for complete details:

SBC1586 (right)
Headless Low-Power Pentium PC/104 board with CompactFlash, 4 COM ports, USB and Ethernet. Starting under $500.

SBC1495 (middle)
PC/104 OEM 486/586 with VGA, COM, Ethernet, and CompactFlash with DOS, Windows, Linux OS.

SBC2590 (left)
EBX Pentium Controller with VGA, COM, Ethernet, A/D, D/A and your choice of Operating System.

preprocessed sensor input data through high-speed A/D boards or front panel I/O connectors such as FPDP. The initial stage of signal processing was formerly handled by ASICs, but it is now more commonly handled by an FPGA-based board on which a customer adds their own proprietary algorithms to the FPGA bit stream. The data is then pushed to a microprocessor-based multi-computer board such as one populated by multiple PowerPCs.

Another factor that makes radar processing in parallel systems especially complex is the need for "all-to-all communication" at certain stages of processing. During processing, radar data must undergo a "corner turn," or a distributed matrix transpose. In a distributed corner turn every processor in the system has a piece of data that every other processor needs for the next stage of processing. In today's VMEbus-based radar systems every processor in the matrix must wait its turn to get a chance to move its data across the bus. Not only does this stall the processing of the current block of data, but it also stalls the incoming data.

Switched fabrics directly address this problem by parallelizing the all-to-all data movement, thus minimizing both the processing stalls and the interruption to the input data streams. Switched fabrics were first introduced into the VME world in the mid-1990s and have grown significantly larger and faster in the intervening 10 years. StarFabric is one serial switched fabric currently used to connect embedded multicomputing products using existing VME64x backplanes. However, physical limits in the legacy VME64x connectors prevent further growth.

Prior to the development of VITA 46, radar systems had begun to confront fundamental performance limits. The two most serious limits being the amount of data that the VME signaling pins could sustain, and the amount of power that could be dissipated per board slot. These are both problems VITA 46 was designed to address. While the number of operations performed on incoming radar data has grown to a certain extent, the amount of incoming raw data has grown tremendously.

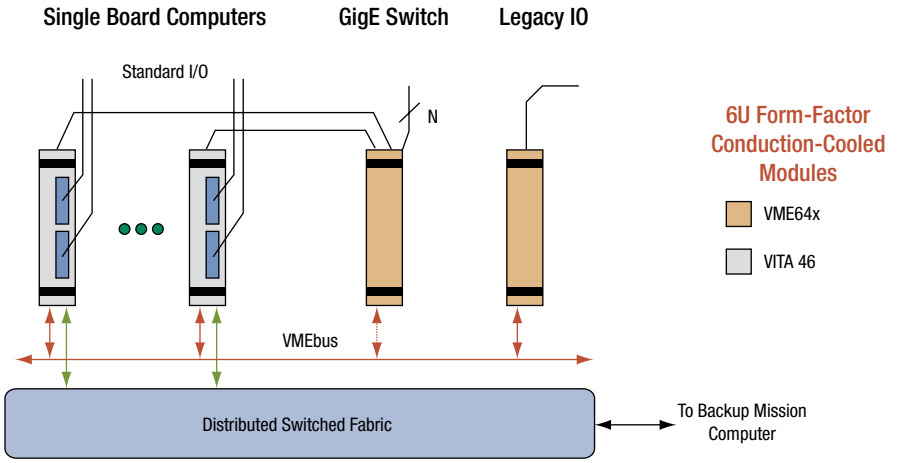

Figure 2 Mission computer—High density of standard I/O, including high-speed serial interconnects (storage), legacy IO (discretes, A/D & D/A) and fabric interconnects including to redundant MC.

Radar data pipes have grown much bigger because the higher the frequency at which data is collected the better the resulting imagery. In addition, the increased availability of large quantities of data is driving even more demanding radar system applications such as moving target identification (MTI), Multi-Hypothesis Tracking (MHT) and Identify Friend-or-Foe (IFF). What's more, there is greater use and deployment of radar sensors. The increased use of UAVs, for example, has increased the amount of data to be processed, while simultaneously, because of stringent constraints on payload size and weight, driving demand for smaller and smaller packages.

As the amount of data has increased, so has the number of processors needed to process the data. This increases the amount of power consumed and the amount of heat generated by a radar system. Unfortunately, the demand for processing power threatens to exceed the ability to cool the components in these systems. This has led some leading vendors to attempt to address the urgent need for higher processor densities with proprietary standards.

However, the industry clearly desires an open standard approach. Smaller radar systems that don't feature such significant amounts of I/O may continue to follow the VME64x path—and for these systems VME or serial switched fabrics on VITA 41 may remain attractive. However, systems that require very large or multiple input data pipes and that use lots of interprocessor communications are expected to migrate to VITA 46.

In the final stage of processing, a radar system sends its data out for operator exploitation. The operator then uses the data, for example, to command a sensor to point in a certain direction or to zoom in on a target of interest. In many cases the operator terminal will be an existing data exploitation terminal with proprietary protocols or interfaces built with legacy SBCs or with custom I/O. With VITA 64, system designers get the best of both worlds, because the new standard provides support for legacy VME64x electrical signals and can utilize a hybrid backplane to house both VITA 64 and VME64x cards. This enables system designers to take advantage of the performance of their distributed fabric of choice while protecting their existing hardware investment.

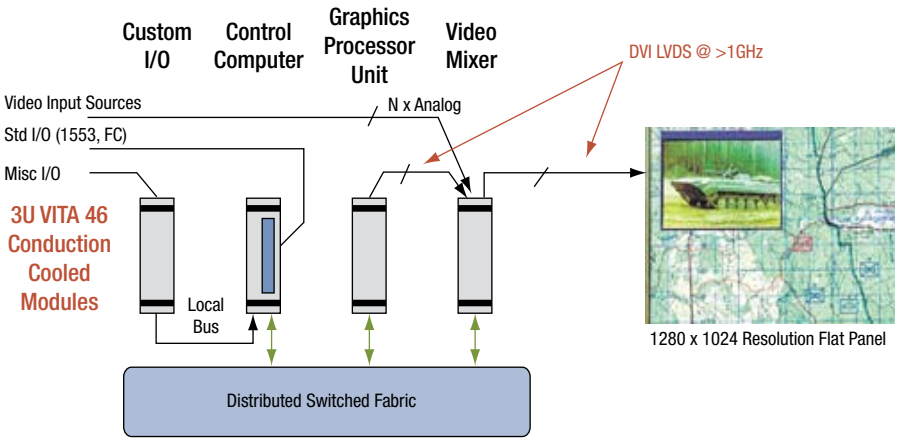

Figure 3 SFF Smart Display—A small form-factor system requires high-bandwidth interconnect for digital video distribution and image manipulation.

Attribute	VME64x	VITA 46
Standard bandwidth	VME: 320MB/s	VME: 320MB/s 32 serial pairs for 10 GBps @ 3.125 GBaud 30 GBps @ 10 GBaud
Switch Fabric	n.a.	Mesh or Central
Faceplate user I/O	Yes	Yes
Backplane user I/O	205 pins	272 pins + 64 pins for fabric
User I/O for 3U system	0	J2 72 pins
Backplane I/O Bandwidth	205 @ 1 GBaud	160 pairs @ 10 GBaud
Defined differential X/PMC I/O mapping	No	Yes
Existing VME64x cards forward compatible	Yes	Yes Use Hybrid backplane
Slot pitch	0.8”	0.8”
Available Power	5v: 90W 3.3v: 66W	5v: 80W 12v: 384W or 48v : 768W
Cooling	Air, conduction	Air, conduction

Table 1 Feature comparison between VME64x and VITA 46

Another advantage of VITA 46 for radar applications is the large number of I/O pins it provides the user in comparison to VME64x. The availability of numerous I/O pins enables system designers to add Gigabit Ethernet off the backplane on VITA 46 and to support multiple ports of rich fabric and the legacy VME electrical signals, as well as provide many additional pins for user-defined general-purpose signals. The abundance of I/O helps to ensure that data flows will keep up to speed, so that no radar image or data will be lost, and no target goes unseen. One more advantage is that control data can be easily split from processing data and delivered throughout the system network via Gigabit Ethernet. This helps eliminate any chance of the high-priority processing data being delayed by lower-priority control data (Figure 1).

Because airborne radar systems commonly have size and weight constraints, it’s expected that some of the early adopters of VITA 46 technology will range from tactical fighters and UAVs to large radar platforms such as JSTARS, AWACS and E2C.

Mission Computers

VITA 46 will also play an important role in the improvement of mission computers. Mission computers provide the intelligent processing needed to handle high densities of standard I/O such as 1553, ARINC, SCSI and RS-422, as well as network fabrics such as Ethernet or Gigabit Ethernet. The mission computer is used to control the I/O or convert the incoming data into commands for fire command computers. It also integrates, consolidates and interprets the data for the user. In addition to handling legacy I/O types, including discrete I/O and A/D and D/A, mission computers are also frequently tasked to control high-speed serial interconnects for data storage and fabric interconnects for a redundant mission computer.

The more data delivered to a mission computer, the better the quality of data available to the system, which results in better and more information that can be analyzed and displayed

The use of switched serial fabrics also opens the door for more distributed video processing solutions by enabling the separation of the video capture and processing functions from the synthetic graphics generation and display functions.

graphically. As with radar processing systems, the mission computer is continually being tasked to handle larger amount of incoming data. This is driven in large part by the proliferation of sensors deployed on a given platform. Unlike the radar processor system though, the mission computer usually has either a single or a small number of CPUs.

However, mission computers are also taking on larger workloads these days thanks to higher resolution cameras, more data sources and high-speed communications. In today's networked battlefield, the mission computer might need to link its database up to the main computer at the central command and control site. In addition to the burden of command and control and security, networked communication also adds the processing burden of data encryption. All of these trends increase the need for faster data rates. As a result, there is a burgeoning demand for higher levels of intelligence, faster I/O and external networking on mission computers.

Another trend driving the need for faster I/O on mission computers is the move away from dedicated platforms, where each displays data from a single sensor or sensor type. The trend today is toward the use of integrated sensors that incorporate data from multiple sensors on a single console resulting in sensor fusion for the war fighter. Increased use of processor-hungry formats such as FLIR and InfraRed are also putting pressure on today's mission computers. At the same time, software is becoming more modular and complex.

As a result, vehicles such as battle tanks and UAVs, which might formerly have had 1553 sensor data coming in, and 1553 and more RS-422 data out, are being enhanced with 1394B, Gigabit Ethernet and FPDP to handle the higher speed data of today's applications. In the past, higher compute requirements had been typically addressed through the use of proprietary secondary buses like RaceWay, SkyChannel, Myrinet and StarFabric, which provide no standard way of connecting multi-vendor components together. VITA 46, in comparison, delivers a standard architecture that eliminates the need for secondary data buses while preserving the user's investment in VME64x hardware and software (Figure 2).

Figure 4 Example of a VITA 46 single board computer from Curtiss-Wright, which features high-speed MultiGig RT interface and high-speed XMC connectors.

Small Form-Factor Smart Displays

The 3U variant of VITA 46 will also find a waiting niche. Today, space and weight-restricted video capture and graphical display systems designed with high-bandwidth interconnects for digital video distribution and image manipulation are frequently built with smaller boards such as 3U CompactPCI cards. These small form-factor intelligent display systems are finding their way onto vehicles like tanks and helicopters where space and weight are a big

PMC Solutions

Reconfigurable Digital I/O

Asynchronous I/O

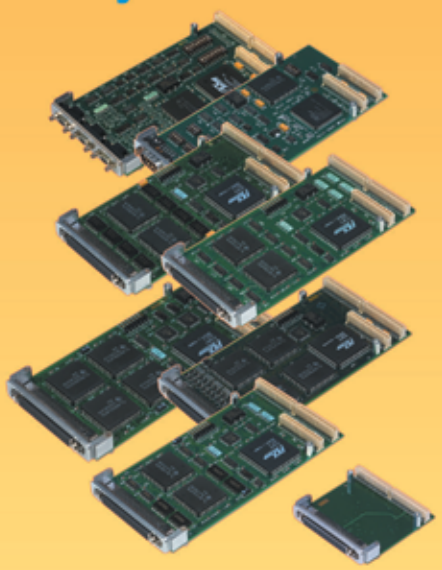

Whenever You Need PMCs ... ... Call Us.

Technobox designs and manufactures custom and off-the-shelf PMCs, including solutions for data acquisition, peripheral connectivity, serial communications, and engineering development.

Our extensive assortment of PMCs includes proven async communication adapters (from dual-port RS232 to 16-port combo boards); versatile, FPGA-based, reconfigurable, digital I/O boards; SCSI (from Fast/wide SE to Ultra 160) and IDE controllers; plus a variety of adapters, extenders, and tools for PMC development and/or integration.

We also offer the industry's only complete x86-based Processor PMC (PrPMC), as well as an ATX-based development platform for PrPMCs.

Processor PMC and PrPMC Platform

SCSI and IDE Controllers

PMC Adapters and Development Tools

For details, visit our web site
www.technobox.com

Technobox, inc.™

PMB 300, 4201 Church Road
Mount Laurel, NJ 08054 USA
Tel 609-267-8988
Fax 609-261-1011

concern. With helicopters, for example, the combined weight of the onboard systems translates into the distance that can be covered with a given amount of fuel, directly impacting which missions can be accomplished.

There is a growing requirement on these platforms to migrate toward the use of higher resolution flat panel displays for mapping, position location and other real-time data display. The move is from today's standard displays (typically up to 1280 x 1024 dpi) to a larger format (typically from 1600 x 1200 and above, including HDTV resolutions) while also enabling system designers to more effectively process and display sensor data by using layers and real-time data updates. This is being hampered by the limits of the connector chain found on traditional CompactPCI systems. The problem is that the greater quantity of data transmitted to the higher resolution displays, and consequently, the higher frequencies and higher signal integrity requirements, cannot always be met by the electrical performance limits of the CompactPCI and PMC connectors.

Customized solutions that do not use COTS can prove to be costly to manufacture and maintain over the lifecycle of the product. The MultiGig connector defined by VITA 46 addresses these electrical limitations. The resulting ability to maintain better signal integrity can also translate to the benefit of driving greater line lengths between the system and the flat panel display. This permits hardware to be placed flexibly in the limited space available in an already dense tank or helicopter interior.

The use of switched serial fabrics also opens the door for more distributed video processing solutions by enabling the separation of the video capture and processing functions from the synthetic graphics generation and display functions. As these two functions no longer need to be so closely coupled, benefits can be realized on both ends of the chain. VITA 46 supports more complex data streams on the input side and more specialized display types on the output side. The increase in backplane data bandwidth provided by the switched fabric also simplifies the addition of video processing stages in the data path (Figure 3). As highlighted by these three application scenarios, VITA 46 already has some hungry customers looking to take advantage of the higher performance and total cost of ownership savings.

VITA 46 products are now starting to appear on the market (Figure 4) and can soon be expected to move in to meet the demands for increased throughput, connectivity and processing power being placed on today's high-end systems.

Curtiss-Wright Controls
Embedded Computing
Leesburg, VA.
(703) 779-7800.
[www.cwcembedded.com]

Removing ATCA's Architected Single Point of Failure

Although a key strength of the AdvancedTCA architecture is its shelf manager, the connections between shelf manager FRUs and Ethernet hub boards have limited overall effectiveness. PICMG's recently adopted ECN 3.0-2.0-001 amends the formal ATCA specification to fix this limitation, while preserving backward compatibility with pre-existing equipment.

by Mark Overgaard
Pigeon Point Systems

The AdvancedTCA (ATCA) architecture is rapidly gaining acceptance among telecom equipment providers and users of that equipment. This architecture's mandatory shelf manager, typically implemented in a dual-redundant configuration by a pair of boards, is a crucial strength of ATCA. Many ATCA shelves implement shelf managers as dedicated field replaceable units (FRUs) connected to the in-shelf Ethernet network. These are also typically implemented in a dual-redundant configuration by a pair of Ethernet hub boards.

Unfortunately, the original AdvancedTCA specification only defines a standard way for each shelf manager to connect to one of these Ethernet hub boards, not to both of them. This limitation can reduce the effectiveness of both the dual-redundant shelf managers and the dual-redundant hub boards, since a failure in just one of each pair may require both pairs to switch over in order to preserve shelf manager connectivity to the in-shelf network. It is more difficult, and much less desirable, to arrange such a coordinated switchover of two separate types of FRUs. In fact, the single shelf manager link to these Ethernet hub boards is the only known architectural single point of failure in ATCA. In contrast, all normal node boards in a shelf have the ability to connect with both hub boards.

Recently, PICMG adopted an Engineering Change Notice (ECN) to PICMG 3.0 R2.0—the formal specification defining the AdvancedTCA architecture—that fixes this limitation, while preserving backward compatibility with pre-existing equipment.

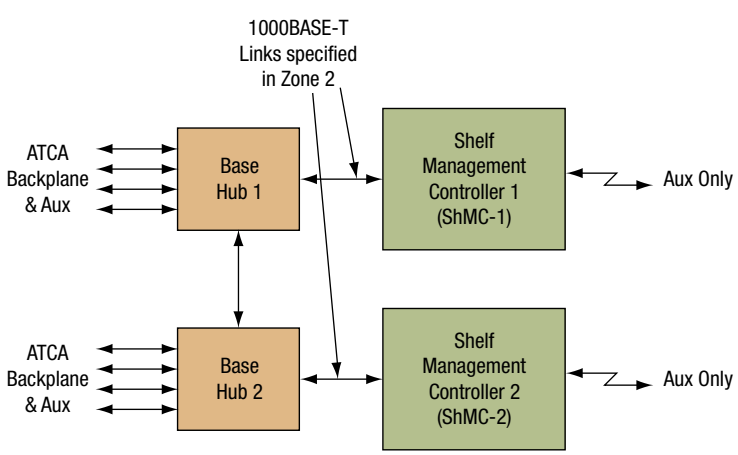

Figure 1 In the original ATCA architecture, each Shelf Management Controller (ShMC) is linked to one hub.

Before ECN: Each ShMC Links to One Hub

In the original AdvancedTCA specification, through revision 2.0 adopted in March 2005, each Ethernet "Base Interface" hub board is connected by a 10/100/1000 Mbit/s link to one of the dedicated shelf manager FRUs, also called Shelf Management Controllers (ShMCs) (Figure 1). This link is provided on the

backplane and reaches the hub boards via their Zone 2 high-speed signal connectors. Other external or auxiliary connections may also be available to each ShMC and to the hub boards. Each 10/100/1000 Mbit/s link requires four differential signal pairs or a total of eight wires.

At the architectural level, there are several consequences of linking each ShMC to one hub. One is that a failure in a ShMC may cause the hub boards to switch over if the interconnects between the two hubs are not used. Another consequence is that a failure in a hub board will definitely cause the ShMCs to switch over. In general, the original architecture reduces the effectiveness of dual redundancy in ShMCs and hub boards due to these effects. AdvancedTCA vendors have developed various ad hoc approaches to remove this single point of failure in their shelf products. However, there has been no generic solution that all vendors can use that preserves interoperability so that a hub board can be installed in any shelf and still implement a solution to these problems.

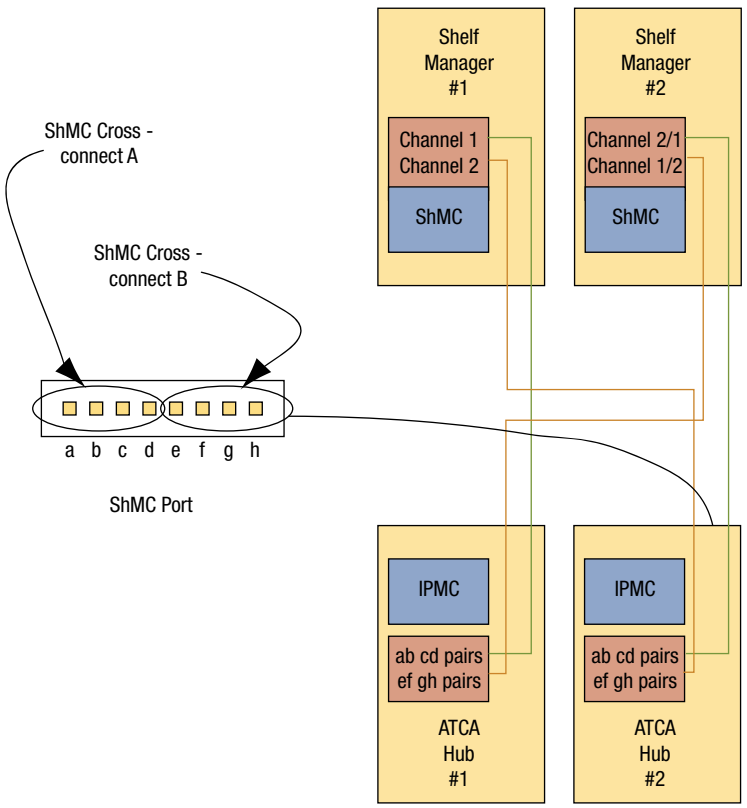

Figure 2 The connections enabled by PICMG's ECN 3.0-2.0-001 include ShMC cross-connects that link both ShMCs to both hubs.

Compatibilities for Hub/Backplane/ShMC Configurations

Config. ID	Connection(s) on Hubs	Backplane Routing	Connection(s) on ShMCs	Notes
1	Dual 10/100	Dual 10/100	Dual 10/100	Normal ShMC cross-connect configuration
2	Dual 10/100	Dual 10/100	10/100/1000	Each ShMC 10/100/1000 channel is connected to two hub 10/100 channels. Hub cross-connect Bs must be disabled.
3	Dual 10/100	Single 10/100/1000	Dual 10/100	No cross-connect; hub cross-connect Bs go to the wrong ShMC. Cross-connect Bs on hubs and ShMCs must be disabled.
4	Dual 10/100	Single 10/100/1000	10/100/1000	Each ShMC 10/100/1000 channel is connected to two hub 10/100 channels. Hub cross-connect Bs must be disabled.
5	10/100/1000	Dual 10/100	Dual 10/100	Each hub 10/100/1000 channel is connected to two ShMC 10/100 channels. ShMC cross-connect Bs must be disabled.
6	10/100/1000	Dual 10/100	10/100/1000	Unpredictable behavior of 10/100/1000 PHYs. Must be avoided.
7	10/100/1000	Single 10/100/1000	Dual 10/100	Each hub 10/100/1000 channel is connected to two ShMC 10/100 channels. ShMC cross-connect Bs must be disabled.
8	10/100/1000	Single 10/100/1000	10/100/1000	10/100/1000 ShMC connections as originally defined in PICMG 3.0 R1.0.

Table 1

The solution for the ShMM-500 was to use a pair of USB links for the redundancy state updates, freeing up both Ethernets for hub communication duty.

After ECN: Each ShMC Links to Both Hubs

The recently adopted ECN 3.0-2.0-001 provides an optional way to remove this single point of failure. It defines dual 2-pair 10/100 Mbit/s ShMC cross-connects that can connect each ShMC to both hubs, as well as map these dual 2-pair links into the single 4-pair connection (the ShMC port) that was originally defined in the AdvancedTCA specification on a hub board for use as an ShMC link. Furthermore, ECN 3.0-2.0-001 adds this option while allowing full backward compatibility for products that conform to the pre-ECN specification. Each 10/100 Mbit/s link requires only two differential signal pairs or four wires.

Figure 2 shows the connections enabled by ECN 3.0-2.0-001, including how the two 2-pair ShMC cross-connects are mapped into a single 4-pair ShMC port on each hub board. When two 2-pair connections—which can each support a 10/100 Mbit/s link—are mapped into a 4-pair connection—which could otherwise support a single 10/100/1000 Mbit/s link—the tradeoff is a reduced data rate for improved availability. However, few, if any, existing shelf manager products support or need a 1000 Mbit/s link to the outside world. For most ATCA users, the availability improvement enabled by the ShMC cross-connects is much more important than the higher data rate.

E-Keying Self-Identifies ShMC Cross-Connect Support

ECN 3.0-2.0-001 also defines new constructs in AdvancedTCA's Electronic Keying facility to allow the three participating shelf component types to self-identify their support for ShMC cross-connects. E-Keying information is represented as records in non-volatile FRU information storage associated with the key participating system components: in this case, the shelf, the dedicated ShMCs and the hub boards. The format of these records is defined in ATCA so that independently implemented system components can self-identify their capabilities, enabling automatic interoperability adjustments.

The shelf FRU information stored with the backplane can show that two ShMC cross-connects are routed on the backplane between each ShMC site and the two hubs. For a legacy shelf, this information indicates just one implemented link from each ShMC site to its corresponding hub. The board FRU information stored on each dedicated ShMC FRU reveals whether they can connect to both of the hubs, if the hubs support that, or to just one of the hubs. The board FRU information stored on the hub boards shows which of the following three possibilities is supported: dual 10/100 Mbit/s ShMC cross-connects, a single 10/100/1000 Mbit/s ShMC link or either of the first two options, dynamically selected when the hub is configured via E-Keying by the shelf manager.

When the shelf comes on, the active shelf manager retrieves the FRU information from each of the above components and enables the appropriate links, based on their capabilities. For instance, the hub boards may be capable of either two 10/100 Mbit/s ShMC links, or a single 10/100/1000 Mbit/s link. Depending on the capabilities of the shelf and the ShMCs, the shelf manager enables one of these two configurations on the hubs.

The Table displays the different configurations of the three critical shelf ingredients for ShMC cross-connects, along with the compatibility properties for each configuration. ShMC cross-connect B is the extra link from a ShMC to the opposite hub that is new with the ECN. In some of the configurations, such as numbers 2 through 5, that link must be disabled by E-Keying on the ShMCs, the hubs or both. Configuration 6 results in unpredictable behavior of generic Ethernet-compliant 10/100/1000 Mbit/s PHYs and must be avoided by system integrators.

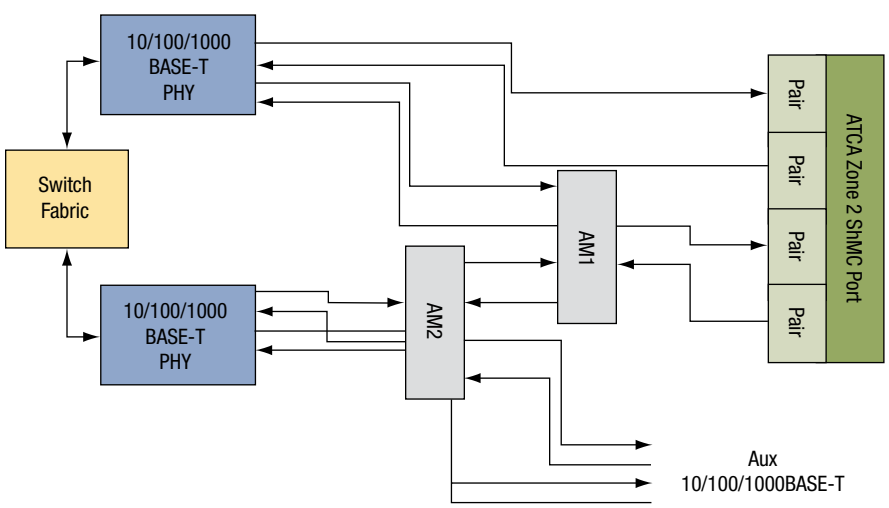

Figure 3 Implementation of a dynamically configurable hub.

Dynamically Supporting One or Two ShMC Links on a Hub

Given the capability descriptions in the E-Keying records, the shelf manager can notify a hub whether to configure

itself for two 10/100 Mbit/s links, or for a single 10/100/1000 Mbit/s link. Such configurability is an optional feature for hub boards that are compliant with ECN 3.0-2.0-001.

Figure 3 shows how such configurability could be implemented in a hub. On the right side of the figure is the 4-pair ShMC port on the backplane connector of the hub board. On the left side are two 10/100/1000 Mbit/s PHY interfaces. Depending on how the analog multiplexers AM 1 and AM 2 are set, one of two possible connections is made between the switch fabric of the hub and the ShMC port. In the first, the top 10/100/1000 Mbit/s PHY is connected via four pairs to the ShMC port. This is the normal pre-ECN 3.0-2.0-001 configuration. In the second, each of the two PHYs is connected with a 2-pair (10/100 Mbit/s) connection to the ShMC port. This is the normal ShMC cross-connect-enabled configuration.

With the implementation shown, the hub can be configured by the shelf manager to support ShMC cross-connects or not, depending on the capabilities of the backplane and dedicated ShMCs.

Providing ShMC Cross-Connect Support

Pigeon Point Systems instigated the PICMG initiative for ShMC cross-connect support in AdvancedTCA and chaired the PICMG subcommittee that developed ECN 3.0-2.0-001.

Prior to the new ECN, some shelf manager platforms, such as the company's popular IPM Sentry ShMM-500 Shelf Management Mezzanine, supported just two Ethernet interfaces: one was available for communication with one of the hubs and the other was used to communicate redundancy state updates between active/backup shelf manager pairs. For such shelf manager products, the advent of the ECN raised the question of how two Ethernet interfaces could be made available for hub connections.

This question was especially challenging for the compact, 67.60 mm x 50.80 mm SO-DIMM-sized ShMM-500, for example, because it would be difficult to add a third Ethernet interface while still maintaining the board's compact size. The solution for the ShMM-500 was to use a pair of USB links for the redundancy state updates, freeing up both Ethernets for hub communication duty. Serendipitously, the four signals needed for such USB links could easily be routed across the same backplane traces that had carried the previous Ethernet-based redundancy state updates.

This approach enabled the quick availability of ShMC cross-connect support within one month after the enabling ECN was adopted by PICMG, allowing users to eliminate the architected single point of failure in AdvancedTCA systems.

Pigeon Point Systems
Scotts Valley, CA.
(831) 438-4709.
[www.pigeonpoint.com].

Intelligent Power Management in High-Availability Applications

High-availability applications such as communications systems require the power distribution architecture and intelligent power management. Successfully implementing both requires careful consideration of power architecture issues, including a power budget analysis.

by Tony Romero
Performance Technologies

Over the past five years, CompactPCI has gained wide acceptance as the architecture of choice for equipment manufacturers developing high-availability embedded systems. Its powerful, standards-based computing solutions, hot-swap functionality and high-availability capabilities make it ideal for applications such as communication. AdvancedTCA (ATCA) is also emerging as a standards-based solution in the communications market. For both form-factors, intelligent power management and the power distribution architecture are critical to the success of high service availability for the complete system. System designers considering the use of advanced managed platforms in their applications must consider several power architecture issues.

Intelligent Power Management

CompactPCI power supplies can be fickle components in high-availability systems. Since today's embedded systems demand higher performance and compute power, power supplies have been stepped up to deliver twice as much power as they did just a few years ago. However, they must deliver that increased power within the same form-factor. These denser power supplies require ample cooling and management, hence the need for intelligent power management.

In the past, power supplies offered simple management information. The two

Application Commands	Sensor/Event Commands	FRU Commands	Firmware Commands	OEM Commands
- Get Device ID	- Set Sensor Hysteresis	- Get FRU Inventory Area	- Start Firmware Update	- Get Device Status
- Broadcast Get Device ID	- Get Sensor Hysteresis	- Read FRU Data	- Erase Flash	- Set Fault LED State
- Cold Reset	- Set Sensor Thresholds	- Write FRU Data	- Program Flash	- Get Fault LED State
- Warm Reset	- Get Sensor Thresholds		- Exit Firmware Update	- Set Inhibit State
- Get Self Test Results	- Get Sensor Event Status			- Get Inhibit State
	- Get Sensor Reading			
	- Set Event Enables			
	- Get Event Enables			
	- Set Event Receiver			
	- Get Event Receiver			
	- Get Device SDR Information			
	- Get Device SDR			
	- Reserve Device Repository			

Figure 1 IPMI command set for typical IPMI-based power supplies.

EBX

PC/104

EPIC

STD 32

Your OEM product can become extinct long before its time, unless you have a long-term commitment from your SBC vendor. VersaLogic guarantees product availability for at least five years from its introduction date.† That saves you from tossing future sales into the jaws of competitors because you suddenly cannot get your hands on the boards you need, when you need them.

VersaLogic's Extended Life-Cycle Policy is just one of the services that helped us earn the prestigious VDC Platinum Vendor award. From our close ties with key vendors to our pre-design component studies, we make it our business to extend your product's life. So sink your teeth into the details at **www.VersaLogic.com/pv13**. Start a long, long-term relationship by calling us at **(800) 824-3163.**

Before Kristin Koons joined VersaLogic's sales support team, she guided tours at a natural history museum. She really knows how to keep products off the endangered species list.

VDC Platinum Vendor

Platinum Vendor award by Venture Development Corporation 2002 and 2003

VERSALOGIC CORPORATION

(800) 824-3163
www.VersaLogic.com/pv13

Figure 2 The CPC6314 from Performance Technologies is a 325W DC power supply that combines high power capacity and industry standard IPMI management capability with monitoring and reporting.

most typical management signals are the Degrade signal, which provides warning when power supply temperature is within 20°C of derating point, and the Power Fail signal, which indicates any output below 90% and/or a low input <36 VDC.

Recently, power supplies have begun implementing the industry standard Intelligent Platform Management Interface (IPMI). IPMI offers a comprehensive set of information to fully manage and predict failures for each power supply. It allows subcomponents from multiple vendors to be monitored by a system's single shelf manager or redundant shelf managers.

The Intelligent Platform Management Bus (IPMB) is a logical bus that was specified to use the I^2C control bus as the physical interface. Power supplies implementing IPMI contain a satellite IPMI/(IPMB) controller. This controller provides an interface to sensors and attaches to the IPMB/I^2C bus as the physical interface.

It processes IPMI commands communicated via the IPMB bus protocol. IPMI is a request-response protocol: the shelf manager, also considered the master or baseboard management controller, issues a request message to an intelligent power supply, also considered the slave or satellite management controller. The supply then responds with a separate response message. Request messages and response messages are transmitted on the bus using I^2C master write transfers. The sensor data record (SDR) is the sensor information stored on the power supply.

Many shelf management modules provide an in-band and/or out-of-band Ethernet interface so that the shelf can be remotely accessed with a common interface by operations, administration and maintenance (OA&M) managers. The shelf manager then communicates to all intelligent components via IPMI.

A power supply that supports IPMI provides five distinct classes of commands (Figure 1). The main commands are: (1) application commands, (2) sensor/event commands, (3) field replaceable unit (FRU) commands, (4) firmware commands and (5) OEM commands.

Application commands initiate operation with the controller, including controller resets and enabling self-tests. The device ID field allows controller-specific software to identify the unique functionality provided by a particular controller. Since multiple power supplies with the same device ID can operate in a system, power supplies can be standardized, yet, at the same time, maintain different geographic locations and FRU information. The Cold Reset command makes the controller reset, while the Warm Reset command makes it start at the beginning of the program, causing the initialization and startup function to be executed.

Sensor/event commands provide the information read from the power supply's sensor data record. Typical SDR information for a power supply includes temperature, voltage and current. The Set Sensor Threshold command lets the user establish different threshold levels for each sensor: minor, major and critical. Each threshold on each sensor can be enabled to generate event messages if that threshold is crossed. For controllers that support the generation of event messages, the Set Event Receiver and Get Event Receiver commands must be implemented. When a significant event occurs, such as a power supply failure, that event must be communicated to the shelf manager as soon as possible, rather than the shelf manager being required to wait to respond to a command. In this case, an event message is initiated by the satellite controller.

The FRU commands provide asset information about the power supply. Asset information is critical to high-availability applications in many ways. The power supply FRU stores the version of firmware it is running. This firmware can be upgraded remotely if required. The FRU information about the power supply—such as part number, serial number, asset tag and date of manufacture—ensures that the technician removes the correct power supply and replaces it appropriately. In addition, knowing which specific power supply is installed can also drive routine FRU replacements, since some OA&M managers proactively replace a component before it fails based on its Failure in Time analysis.

Firmware commands can update the power supply effectively and remotely. These commands can also be used to erase flash memory space occupied by the program code being updated and to program the flash.

OEM commands allow users to change the state of the power supply either locally or remotely. This includes commands that monitor and override the state of the Power OK LED and the power supply shutdown output, a device status command, device-implemented FRU-type commands and device set-up commands.

The ATCA specification eliminated independent power supplies from the platform, but ATCA power entry modules (PEMs) also include IPMI controllers. PEMs take the -48V input from telecommunication facilities and distribute it to various slots and supporting components in the platform, such as fans. Each board must convert the -48V into the voltage levels needed. Typical ATCA platforms support two redundant PEMs. The shelf manager can monitor the voltage, current and thermal sensors on the PEMs. If one PEM fails, the shelf manager will be notified so the failed PEM can be replaced.

Some power supplies combine high power capacity and standards-based IPMI management capability with monitoring and reporting capability (Figure 2). These critical elements, along with redundancy and hot swap, give embedded applications the performance and management capabilities in power supplies they need to achieve high availability.

Power Budgeting

When dealing with high-performance applications in a CompactPCI platform, it is important to develop a power budget analysis to ensure that there is sufficient power for the complete system configuration. High-performance applications tax both the power and cooling architectures supported in the platform. Power budgeting can also provide a rough estimate of how much average power can be delivered to each platform slot.

To properly budget power for 12U platforms, it is best to develop a spreadsheet to account for all components that draw power. CompactPCI defines four independent and limited voltage rails that deliver power to the platform: 5V, 3.3V, 12V and -12V. It is important to analyze power budgeting at each voltage rail, because one specific configuration of boards in a chassis may not tax the 5V rail but could significantly tax the 3.3V rail, for example. Each voltage rail is independent and does not share current. Thus, when one rail's limit has been reached, the limit of the number of boards that can be integrated into the chassis using that specific voltage has also been reached.

The power budget analysis begins with a list of the maximum current rating for each component and board to be integrated into the platform. For each unique board, the designer should write down its maximum current draw in amps per voltage rail. Using the maximum power draw for each component is necessary to calculate a worst-case scenario.

The next step is to list all boards and components to be configured in the platform, including quantity. Figure 3 shows a sample power budget analysis spreadsheet in table form. Each row represents the maximum current rating for the total number of units configured in the platform. This is broken out into the four voltage rails. At the bottom of the table, the total current for each voltage rail is added up to determine if the power supplies provide ample power. If the analysis shows a deficit in the current available for one of the power rails, the designer has several options. These may include removing one payload from the system, using alternate boards that consume less power, or looking for power supplies that can deliver more power.

Power Budget Analysis for Configured Platform
Enter your values into the white boxes. The gray boxes are results and calulations

Part #	Units per platform	+5V (A)	-3.3V (A)	+12V (A)	-12V (A)	Watts
Wire Harness	2	6.00	12.80	0.83	0.00	82.20
ZT 5061	3	0.00	0.00	9.00	0.00	108.00
CPC6600	2	13.52	12.74	0.02	0.02	110.22
CPC7301A	2	6.00	0.00	0.00	0.00	30.00
CPC5505	4	11.28	36.00	0.00	0.00	175.20
CPC388	3	0.00	15.30	0.00	0.00	50.49
CPC358	5	2.50	29.00	0.05	0.05	109.40
0U812	4	24.00	56.00	4.00	0.00	352.80
CPC6314	8	0.48	0.00	0.00	0.00	2.40
N/A	0	0.00	0.00	0.00	0.00	0.00
N/A	0	0.00	0.00	0.00	0.00	0.00
Rear Transition Boards						
RTM6600	4	0.00	0.00	0.00	0.00	0.00
RTM4820A	1	0.20	0.00	0.00	0.00	0.98
ZT 4807	4	0.00	0.80	0.00	0.00	2.64
RTM388	3	0.00	0.00	0.00	0.00	0.00
RTM358	5	0.00	0.00	0.00	0.00	0.00
N/A	0	0.00	0.00	0.00	0.00	0.00
N/A	0	0.00	0.00	0.00	0.00	0.00
N/A	0	0.00	0.00	0.00	0.00	0.00

Total Current Draw (A)		63.98	162.64	13.90	0.07	
Total Draw (W)						1024.32
Total Current Available		120.0	160.0	20.0	4.0	
Delta		56.0	(2.6)	6.1	3.9	

Figure 3 In this example of a power budget analysis, there is a deficit in the current available for the 3.3V rail. The configured boards demand a total of 162.64A, but the power supplies are only delivering 160A.

Another question to keep in mind in a CompactPCI or AdvancedTCA deployment is whether the facility provides enough power to the rack or frame. The power budget analysis shown in Figure 3 assumes that the facility can provide sufficient power to the power supplies for CompactPCI and to the PEMs for AdvancedTCA.

Increasingly, designers are realizing the benefits of higher levels of availability. Not only does this provide them with a leg up on the competition, but it also reduces total cost of ownership during the life of deployment. System management is an important element in this equation, and power supplies are no exception. They need to be managed intelligently like any other high-availability board or component in the platform. With never-ending performance enhancements, it is important to keep in mind the analysis necessary to ensure power supplies can power the complete configuration. ◢

Performance Technologies.
Rochester, NY.
(585) 256-0200.
[www.pt.com].

ALL APPLICATIONS LEAD TO...

Compact PCI®

PCI CONVENTIONAL PCI-X

ONE STOP SYSTEMS

PCI EXPRESS®

MAX EXPRESS

ISO 9001:2000

PCI SIG

ONE STOP SYSTEMS
Corporate Headquarters
2235 Enterprise, Suite 110
Escondido, CA 92029
Tel (760) 745-9883
Fax (760) 745-9824

Call today.

(877) 438-2724

www.onestopsystems.com

Executive Interview

RTC Interviews Tom Quinly, President of Curtiss-Wright Controls Embedded Computing

***RTC*: VMEbus, now stretching to almost a quarter century of service in the military and elsewhere, continues to be the leading bus architecture in the military space. To date, VME has by and large not been displaced and, in fact, most of the current VME systems deployed use the traditional VME64. How long can the aging architecture continue to dominate the standards-based board market in the military? Why?**

Quinly: Without question, the longevity of the VMEbus, and its continued popularity, results from the progressive steps of improvement that have enhanced the standard over the years. For example, the evolution from 32-bit to 64-bit, and the development of ruggedization and cooling variants. The way that VITA and the embedded community have nurtured and improved the standard has led to stability, and the military and other markets value stability. Our customers rely on these solid standards.

At this quarter-century mark we're about to see a new step change in the evolution of the VMEbus: VITA 46 will deliver new high-bandwidth switched serial fabric technology and features like 2-Level maintenance, coupled with the stability that the market demands for the next quarter century. That's not to say that VME64 will go away. For our customers, when it's time to start cutting new platforms over, it's done platform by platform when the right economic reasons are in place. Our customers don't change for change's sake. They design a box and live with it for a long time. When customers decide to change, they want a platform that can take them through the next 25 years, and VITA 46 is the major step function for the next quarter century. In terms of time, we think that the upgrade from classic VME and VME64 to VITA 46, for many customers, will be a 5 to 10-year crossover period. And for some, VME64 will remain an affordable, cost-effective solution. But when they need to make the move because of features or to support their application, we'll be there for them.

At Curtiss-Wright we're committed as a company to support our customers who deploy VME64 platforms. Thanks to the size of our company, we have the financial muscle and talent pool to support both legacy VME and VITA 46, so our customers can do the crossover when the time is right and their application requires it. We'll continue to provide the latest technology in both VME and VITA 46, and we will work to make sure that these products have a common set of features that serve our customers, what we call "COTS Continuum," so that when they do change over to VITA 46 they are assured of the advantages of working with the same vendor, same API and the same ecosystem. As far as the VMEbus is concerned, we're confident that we'll be designing new VME cards for another 5 years and building them for another 15. We have assembled the industry's best lifecycle maintenance program and we have longevity and supply commitments on all VME variants. But the dominant story for the next 25 years is VITA 46.

***RTC*: The VSO has been working diligently on several new specifications including VITA 41 and VITA 46. At the present time, Curtiss-Wright as well as many other companies, have put a great deal of resources into the specification processes. Do you believe one, or some combination of the two (VITA 41 and VITA 46) will replace the traditional workhorse VME? In that context, VITA 46, as the first VSO specification to totally separate from the traditional VME spec, does not include any VME pin-outs at all. However, with its high-speed, high-density connectors, it's expected to be ready for some new applications as they emerge. Do you, at the present time, have any feeling for the acceptance of VME 46 in advanced**

Copenhagen
San Jose
Albuquerque
Tucson
Phoenix
Boulder
Salt Lake City
Milan
Huntsville
Melbourne
Rome
Minneapolis
Chicago
Lucerne
Madrid
Barcelona
Boston
Greenbelt
Houston
Dallas
Indianapolis
St. Louis
Oslo
Toronto
Montreal
Gothenburg
Ottawa
Helsinki
Long Beach
San Diego
Pittsburgh
Morristown
Washington DC
Eindhoven
Atlanta
Raleigh
Portland
Seattle
Vancouver

You're Invited

Atlanta, GA
November 15
Raleigh, NC
November 17

Eindhoven, Netherlands
November 2

Portland, OR
December 1
Seattle, WA
December 6
Vancouver, BC
December 8

Real-Time & Embedded computing conference

Enter a World of Embedded Computing Solutions

Register now at www.rtecc.com

Sessions

Through open-door technical sessions especially for those developing computer systems and time-critical applications. Get ahead with sessions on Embedded Linux, VME, PCI Express, ATCA, DSP, FPGA, Java, RTOS, SwitchFabric Interconnects, Windows, Wireless Connectivity, and much more.

Exhibit Hall

Arranged in a unique setting to talk face-to-face with technical experts. Table-top exhibits make it easy to compare technologies, ask probing questions and discover insights that will make a big difference in your embedded computing world.

Sponsors

AAEON · Accelerated Technology · ENEA · Green Hills Software · LynuxWorks · SBS Technologies · WDL Systems · Wind River

military platforms? What platforms and in what applications do you believe it would be most suited?

Quinly: Both VITA 46 and traditional VME will survive and co-exist for many years. But what the market and our customers are telling us is that the breadth and depth of the heart of the industry will move, over time, to VITA 46. Military customers are pragmatic and favor the long-term view; they don't like stepping stones and half measures. They need all of the advantages that VITA 46 delivers. While we see VITA 46 as the ultimate evolution path, customers will have different paces for getting there. And, there may be market segments that don't put as much value on all of the VITA 46 features. For applications that aren't purely conduction-cooled, a market segment where VITA 46 will clearly dominate, other point solutions such as air-cooled VITA 41 will have a place, frequently as a bridging strategy. And Curtiss-Wright will support some VITA 41 point solutions for customers with special requirements. But the bulk of the industry is setting their product evolution right now and we're seeing the mainstream heading toward VITA 46.

For our customers, exact connector compatibility between VME and VITA 46 is less important than board to board compatibility, and we've given them that in the spec. The VITA 46.1 specification defines VMEbus electrical signals on particular pins of the VITA 46 backplane. With heterogeneous backplanes customers can transition and build systems with which they can still leverage their investments in VME64x because the VME electrical spec is still carried across. Since providing identical pin-outs came at too high of a performance tradeoff, the VITA 46 Working Group took the next best approach, which is to design systems that allow interoperability between VME 64x and VITA 46. And when we canvass our customer base, that's what we've found people really care about. This approach lets them solve their future problems with VITA 46, but if they still want to use some of the existing products that they've invested in, especially their more costly specialty products, this approach allows them to do exactly that.

It's clear that for many of today's emerging applications and programs, such as FCS, where you're going to be doing much higher density computing in ever smaller platforms, VITA 46 is the way to deliver the high-speed fabrics and extensive fabric interconnectivity that are required.

***RTC*: CompactPCI, both in its 3U and 6U form-factors, has been gaining increased acceptance in the military recently. What are the advantages—if any— that they bring over traditional VME?**

Quinly: We see 3U CompactPCI as a winner in its niche. It addresses a real need in the marketplace, which is to provide a standard form-factor for smaller size systems. And thanks to the improved processing capability of chips and increased memory densities, it's now viable to get high-functionality cards that formerly required a whole 6U card. For this class of applications, 3U cPCI offers a great power/weight/performance equation. But for newer generation processors, CompactPCI faces the same roadblock that VME does. Its connectors simply can't carry the signaling speed required to support the future technologies. We expect that 3U cPCI will carry on, but will eventually also morph into the already defined 3U variant of VITA 46.

***RTC*: PICMG just approved its CompactPCI Express specification. It has been defined in a variety of form-factors. Do you believe the new specification will have a place in the market given the time lead that VITA 41 has? Are there performance issues that you can envision that would make one more attractive than the other in particular applications?**

Quinly: We don't see demand for CompactPCI Express (cPCIe) as a standard in the military market. It is aimed at smaller systems that can live with PCI Express's constraints for peer-to-peer computing. For larger systems you'll need RapidIO or Advanced Switching, PCI Express by itself won't do the trick. Also, there's no mechanical Compact PCI Express spec other than commercial, and if developed, a rugged implementation would not be able to fully support PMCs.

CONCEPT DEVELOPMENT, INC.

SENIOR ANALYST

This position requires a Senior Analyst with strong system analytical skills, excellent verbal and written communications skills.

The Senior Analyst works closely with business clients, internal Project Management and Application Engineering to analyze requirements, formulate proposals, provide systems architecture and craft implementation strategies. The solution space is complex Field Programmable Gate Array (FPGA) core design, real time embedded firmware or host processor functions. An essential requirement of this position is the ability to clearly articulate, in writing, specifications. These specifications include strategic architecture as well as detailed implementation requirements that can be turned over to the implementation organizations. The Senior Analyst must be able to craft architectures and strategies that are consistent with the constraints imposed by customer's needs, the skill levels of the implementers, the time and budget allotted to the project and feasible technologies. The ability to get to the heart of complex technical issues and establish proper perspectives in an active leadership role is paramount as is the demonstrated ability to conscientiously and effectively work to tight schedules.

As such, the Senior Analyst must have a practical flare and clearly understand the impacts of architecture and design upon cost and schedule. At the same time, the Senior Analysts must be creative and open to novel approaches to problem solving in a multi-discipline environment.

A demonstrable history of experience in the fields of imaging, signal processing, complex algorithm implementation or communications with FPGA, ASIC, DSP or embedded processors is essential.

SKILLS:

- At least ten years experience in complex Systems Development
- Strong system analytical skills
- Experience or familiarity with FPGA or ASIC design
- Experience or familiarity with embedded firmware or software design
- Experience or familiarity Object Oriented Design and UML
- Excellent verbal and written communications skills
- Competency with Microsoft Office Suite, including Project and Visio
- Proficiency in data modeling
- Background in preparing white papers and proposals.
- BS/MS/PhD in Engineering or Physical Science (mandatory).

ANALYST

This position requires an entry level engineer or science major with good analytical skills, and above average verbal and written communications skills.

The Analyst works closely with business clients and Application Engineering to analyze requirements, formulate proposals, provide systems architecture and craft implementation strategies. The solutions developed for customers include complex Field Programmable Gate Array (FPGA) core design, real time embedded firmware or host processor functions. An essential requirement of this position is the ability to clearly approach problems and communicate design concepts. The ability to conscientiously and effectively work to tight schedules is required.

The analyst works with a team of Application Developers to being client requirements to fruition in terms of a real product. These products span industries and technologies including image processing, communication, command and control, medical and aerospace applications.

A demonstrable history of experience or academic course work in the fields of imaging, signal processing, complex algorithm implementation or communications with FPGA, ASIC, DSP or embedded processors is desirable.

SKILLS:

- Good analytical skills
- Experience or course work with FPGA or ASIC design
- Experience or course work with embedded firmware or software design
- Experience or course work in Object Oriented Design and UML
- Above average verbal and written communications skills
- Competency with Microsoft Office Suite, including Excel and Word
- BS/MS/PhD in Engineering or Physical Science (mandatory).

Go to **www.cdvinc.com** or call 949-623-8000.

The VITA 46 standard has means to support interconnection to legacy CompactPCI modules to harmonize the go-forward route to switched fabrics for current users. I see it as a winner for some commercial applications in 3U, but larger apps will likely go ATCA. We also expect that current rugged users of 6U CompactPCI will migrate to VITA 46.

***RTC*: There has been over the past few years a major move toward consolidation within the mid- and high-end of the embedded-computer industry. SBS, Curtiss-Wright, RadiSys, Motorola and most recently Mercury Computer have been aggressive in their consolidation moves. What do you believe are the benefits that consolidation brings the customer? What are the disadvantages?**

Quinly: By bringing together best-of-breed companies, such as Dy 4, VISTA Controls, Synergy Microsystems, Systran, Primagraphics and Peritek, Curtiss-Wright Controls Embedded Computing is able to support the customer from design to production, from mezzanine card, to processor board to subsystem enclosure, with off-the-shelf, semi-custom or full custom services. Even better, our enhanced strength provides significant improvements in customer service, including lifecycle maintenance, which is a key issue in the military market.

It's important though to distinguish between companies that grow by adding various centers of excellence strategically to better support the customer as we believe we have, and those that consolidate simply in an attempt to inorganically add revenue. Our acquisitions have been intelligent and designed to add complementary, synergistic capabilities, from the card all the way up to fully integrated subsystems. This is also where our COTS Continuum initiative plays an important role by creating a standard user experience for development across our product lines. That consolidation results in a coherent, consistent ecosystem that makes a customer's life easier rather than merely creating a collection of individual legacy companies. Keep in mind that in the embedded board market the biggest challenge has always been to convince customers to outsource. Innovative consolidation and the strength of a parent like Curtiss-Wright have made us a bigger and better fit for these potential customers.

***RTC*: Since its inception, the mid- and high-end embedded-computer industry has been a platform for prototyping, beta development and very low-volume production. Do you believe this scenario is changing and that more and more production orders will be going to board and subsystem vendors? If so, what is driving it? If not, what is inhibiting it?**

Quinly: The low-volume production business is where COTS has historically offered the best value proposition. But as the embedded market matures, our ability as an industry in general to take on and deliver to major programs continues to improve. For Curtiss-Wright, as a bigger company, we've increased our manufacturing capability, support capability, our capital and our infrastructure to reliably deliver and service major programs over the twenty years of their lifespan. In the past, low-volume COTS production never translated into big-volume production; these programs would get subsumed in-house because the big companies never expected that COTS companies could be relied on for their production programs.

What we see now is that there are fewer large programs in the traditional sense. They've become less frequent because of the whole paradigm of continuous spirals, faster time-to-markets and fewer big volume runs. These days, a large volume program is more accurately characterized as five small volume production runs of technology refreshes over the program's lifetime. Five years ago, for example, the individual companies that now comprise Curtiss-Wright Controls Embedded Computing wouldn't likely be considered for larger programs. But now, as part of a large legacy company, we've expanded the longevity of our repair capabilities and we're seeing more subsystem and integrated solutions business.

***RTC*: The terms "semi-custom" and "semi-standard" have become almost synonymous with the board and subsystem business. One vendor was quoted as saying, "virtually everything coming out of our shop has some level**

of customization." This is often the case as basic standard designs are adjusted to fit particular requirements. Have the economics of outsourcing changed the structure to allow companies to accept more standardized components? Or, has just the opposite happened whereby more fully custom platforms are developed? Does this differ from the commercial/industrial market to the military market?

Quinly: One notable trend we're seeing is that, in general, military customers are more willing to accept standard products. The big enemy of customization is time-to-market. Related to the time-to-market issue is cost of goods, and today customers just don't see the burden of price on standard products. For example, three years ago you'd expect to see prices starting around $25K for a Quad PowerPC board, today the starting price would be closer to $15K. An SBC that used to cost $17K is now likely to cost $10K. Off-the-shelf, standard products have become more competitive. And the trend is for people to pick a leading off-the-shelf technology. This drives our commitment to offer a leading product portfolio. We put a lot of R&D dollars into maintaining our technology leadership with standard products.

That said, we do have both standard and custom business. We're structured with both a subsystems group and a modular services product group. While we're not positioned as a custom shop, we do see custom business and we structure our business to support it both at the board and subsystems level. However, to be effective as a custom partner requires a portfolio of state-of-the-art standard COTS products and a design infrastructure to support design re-use. It's our wealth of COTS intellectual property that makes custom programs work; standard products facilitate the speed and cost with which you can provide custom offerings. They're linked hand in hand. Because NRE can swamp production builds, COTS designs have a better chance of living all the way to production.

One of the weak areas in our market is that customers often don't consider COTS solutions when they should; when they want semi-custom they should be talking to people like us, because we can offer them great solutions for semi-custom systems that leverage our COTS IP. Many don't think about it, but some smart customers are starting to recognize that opportunity. For example, COTS is increasingly attractive economically for big programs, which have chosen to go with non-standard form-factors, such as the Joint Strike Fighter. The program may be too far down the road to change the computing module form-factor, but it can certainly leverage COTS IP, particularly as it gets ready to go to production and needs a technology refresh. And newer programs, like FCS, are now looking at VITA 46.

***RTC*: Curtiss-Wright Controls addresses both the military and commercial/industrial markets. Can you share with our readers the following: (1) What, outside of the obvious ruggedization issues, are the differences between the two markets? (2) What proportion of CWC's business is devoted to non-military applications?**

Quinly: The main differences are volume and application-driven features. Simply stated, volume drives pricing. And the commercial/industrial markets see completely different volume levels than we see in military programs. The need for high-performance, mission-critical, ruggedization and advanced cooling requirements in the military also drives the use of features that may not be required in commercial/industrial environments. Greater-volume commercial customers expect to see their price point adjusted—they're not typically willing to pay for any legacy or genealogy of ruggedization in their commercial product. When we design a product we choose to make it available in a rugged flavor for the military. To take it backward and offer it competitively in the commercial market is difficult because the product always retains a design legacy of ruggedization.

Increasingly though, time-to-market issues are becoming similar in the military and commercial markets. Another way the military market is becoming more like the commercial market is that our customers are more frequently pushing to have our latest/greatest leading products as soon as possible. Curtiss-Wright is consistently an early adopter and Beta site of new enabling technologies (e.g., processors) to achieve our time-to-market goals. As far as our interest in the commercial market, our strategy is to be very selective in market verticals where we can bring value, such as semiconductor test equipment, simulation, oil and gas.

***RTC*: In general, what has the effect of COTS been over the past few years for your business as well as for the market in general?**

Quinly: COTS has continued to grow in acceptance. We see corporate-level decisions being made by leading systems integrators to leverage COTS instead of continuing down the in-house "make" route. COTS has become the norm, and in the process it's created a growing expanding marketplace, which we've taken advantage of by broadening our capabilities and offerings. Evidence of this maturation is the extent to which COTS has attracted the interest and support of large established parents like Curtiss-Wright, who by investing in it are helping to bring the market and technology to the next level.

***RTC*: The communications/networking infrastructure market is starting to heat up as a result of convergence of voice, data and video and the battle for the living rooms between the traditional communications providers and the cable companies. This has resulted in some early activity on the ATCA and AMC front. Could you comment on the emergence of these technologies from the perspective of your company? Do you see much (or any) interest in the military market for them? Does Curtiss-Wright Controls Embedded Computing plan to participate in the communications/networking market? If so, what platforms and technologies will it embrace?**

Quinly: We have no plans to enter the telecom market; it's just too different from our expertise. We believe that to be successful you have to be #1 or #2 in your market. That said, the trend toward network-centric warfare is driving us to bring leading-edge telecom technologies to the deployed defense market. We already offer Gigabit Ethernet switching products that are field-deployable in module or box-level solutions. So while we don't play in the commercial communications/networking market, we are leveraging advanced solutions into MIL applications.

For example, we deliver intranet-working platforms that tie together various sensor and intelligence resources within a land-based vehicle or on a UAV platform, and we also deliver internet-working platforms that tie these individual platforms together on a networked battlefield to distribute real-time data to the warfighter. Some of these applications are located in backrooms and protected environments, but increasingly, the paradigm of a distributed command and control structure is driving computing platforms directly out to the warfighter. This trend places telecom and network technologies onto moving platforms where shake and vibration survival is critical and much better accommodated by standards such as VITA 46. If there are applications for ATCA and AMC, they won't be in those deployed environments.

Embedded Network Attached 6U Storage Blades Improve System Availability

Embedded systems typically use mass storage as part of the overall system implementation. Today, designers now have a choice between a conventional parallel bus design and a fabric-based network system architecture that provides highly available (HA) systems. A primary reason for designers to consider fabric-based systems is scalability and availability or redundancy.

by Ken Grob
ACT/Technico

In the embedded space, the PICMG 2.16 cPSB specification and the VITA 31.1 standard are two standards for embedded system area networking. The first uses the standard CompactPCI form-factor. The second uses the VMEbus form-factor. These two standards, which make use of the 6U by 160 mm Eurocard form-factor, allow the system designer to build loosely coupled systems, using embedded physical network connections in the system chassis. To complete such designs, network attached storage (NAS) at the blade level is a useful, and much needed, component.

In order to support the construction of a highly available system, a designer will typically eliminate single points of failure and build in redundancy where possible. In order to increase uptime, or availability, the designer must increase the mean time between failure (MTBF), most commonly by eliminating dependencies between components. Designers can achieve an increase in system uptime by:

- Using highly reliable system components
- Providing a redundant infrastructure, including cooling, power sources and power supplies
- Providing redundant networking paths
- Using loosely coupled, redundant processors

Two approaches are typically employed to provide redundant system resources. The first is complete system redundancy achieved by total hardware duplication, where the systems are run in an active/stand-by fashion. Alternatively, a distributed system of nodes can provide for graceful degradation. Typical components in this system include the system chassis, the fabric switch and the single board processor. The PICMG 2.16 cPSB specification allows scaleable systems to incorporate a Dual Star network topology containing N node slots, where N can vary from 1 to 19 slots (Figure 1).

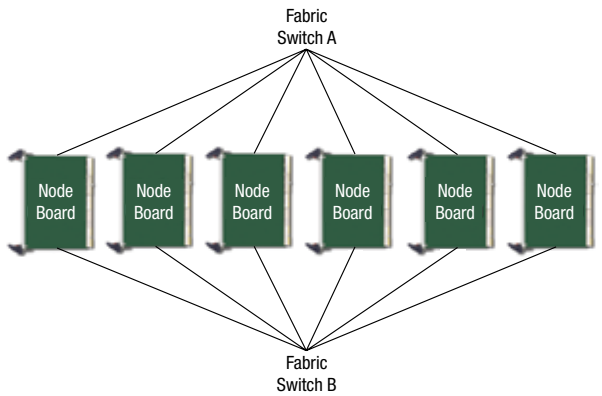

Figure 1 Dual Star network topology.

When reviewing the overall system architecture, the designer must consider storage. Storage is included on a per blade basis for local access. Global storage is typically considered external to the chassis with the interconnect via an adapter interface such as SCSI or Fibre Channel. In analyzing the need for storage, its use must be considered too, which can be broken down into application-specific categories, including operating system boot and load; application image load; data storage and logging of events.

System designers need to consider where the single points of failure are located in order to eliminate them. In a system using multiple single board computers (SBCs), each with its own drive, the boot device would be a single point of failure. Therefore, having disks located on each SBC in the system will introduce single

points of failure. Other issues include sharing data or making data available to multiple processors, where a file system must be shared or be accessible to multiple computers. Placing a disk on each SBC in the system prohibits all of the processors from sharing and conveniently storing the data. Networked storage is more desirable.

In addition, one needs to consider data replication to increase availability of the stored application data. Replication of the file systems becomes an issue of maintaining data coherence across the redundant storage volumes. A scheme is required for duplication. Secondly, it is necessary to have a method of sharing that does not burden the application programmer relative to accessing system storage.

Blade-level NAS can address storage needs for a scaleable system based on PICMG 2.16 or VITA 31.1. The NAS blade is an embedded network storage appliance that connects via the embedded system area network. The NAS blade provides the necessary internal processor, network interface and disk storage in a single slot Eurocard form-factor. NAS provides a client/server model for storage access. At the blade level, the NAS blade is viewed as the storage server, using familiar protocols such as NFS 3.0 and CIFS (Common Internet File System) (Figure 2).

In order to design a highly available storage solution, one must first identify what elements will fail and how to avoid those failures. A designer should therefore consider functional issues related to:

- Replication of data over multiple disks to cover disk failure
- Controller and interface replication to cover controller failure
- Data duplication between storage blades to cover blade failure
- Replication of the network path to cover the failure of a network interface
- Access to the data by the application program

Implementing a high-availability NAS blade requires a set of features that address the above functional issues. These can include multiple disks organized as a RAID set within the blade, transparent duplication of the user data between blades and automatic failover between redundant NAS blades. Equally important is support of redundant network interfaces, heartbeat and failover between blades and a method of user interface and event detection, as well as hot swap of the blade and drive.

Additionally, application issues such as storage capacity, bandwidth and access times must be considered.Capacity is supported by the capacity per drive. Using 2 ½" SATA drives, capacity is now approaching 100 Gbytes per drive. Bandwidth is a function of the network interface, drive and onboard NAS processor. In order to provide local redundancy, the architecture of a NAS blade must address the necessary feature set and include a minimum of two disk drives per blade, where the drives are used in a RAID One (mirrored) configuration. Figure 3 shows such a blade, including hard drive shuttles.

To provide a robust storage implementation, the NAS blades must provide a method for data duplication. Using RAID methods, and logical block duplication, the NAS blade architecture provides for inter- and intra-blade redundancy by replicating user data across drives contained within the blade in a single blade application, and between the blades in a dual blade application. When two blades are used, the data is duplicated four times. In this way, a single drive and an entire NAS blade failure are covered by a redundant resource.

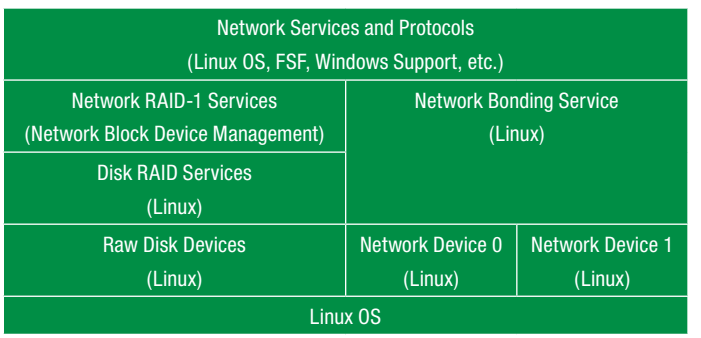

Figure 2 NAS software model.

In order to further increase uptime, front panel hot swap of the backup blade or redundant drive can provide replacement without service interruption by using a drive shuttle with SATA interface. There must be a way to allow the drives to be hot-swappable from within the blade. Software in the blade detects removal and insertion of a drive. When a drive is replaced, its data is automatically updated on the blank drive to restore the redundant media. Complete blade replacement can be addressed

Figure 3 NAS blade showing dual drive shuttles.

	Single Blade: Mirrored Drive	Dual Blade: Intra and Inter Blade Mirror
File System Protocol:	NFS	NFS
Read Performance:	8 to 12 MB/Second	8 to 12 MB/Second
Write Performance:	7 to 10 MB/Second	4 to 5 MB/Second

Table 1 NAS Performance Summary

by removing a blade, making data refresh seamless (Figure 4).

In order to remove a blade, the NAS server function must be transferred to a backup blade. This is accomplished through a heart beat program and a NAS IP alias scheme, where the active network device IP address switches from the primary NAS blade to the secondary blade, as a result of a failure. This action, called failover, occurs when the secondary NAS blade becomes the active server, supplying data to requesting applications and allowing the primary blade to become available for removal or replacement.

Upon replacement of a swapped blade, the data must be rebuilt on the new device. Detection by the NAS blade software of an out-of-sync blade allows it to be rebuilt across the network in the background. While this occurs, the now active secondary blade is still available to serve data until the primary blade is replicated. This function requires capped, in-band network bandwidth. An additional out-of-band path for data replication would further increase the availability and reduce the replication time between NAS blades.

NAS Management and Event Detection

Although an NAS device is designed to be reliable, it still requires event or alarm detection and notification. If an error occurs—for instance when a disk or blade fails—the storage complex must be able to send notification. This can occur through an alarm system that provides status and event information, such as Simple Network Management Protocol (SNMP), where traps are generated and presented to an agent process associated with an SNMP manager. Here, as events occur, the agent process receives the SNMP trap and provides status to the user via a system management interface. Another method is to use a Web-based graphical user interface (GUI), where a window is updated as status information comes in.

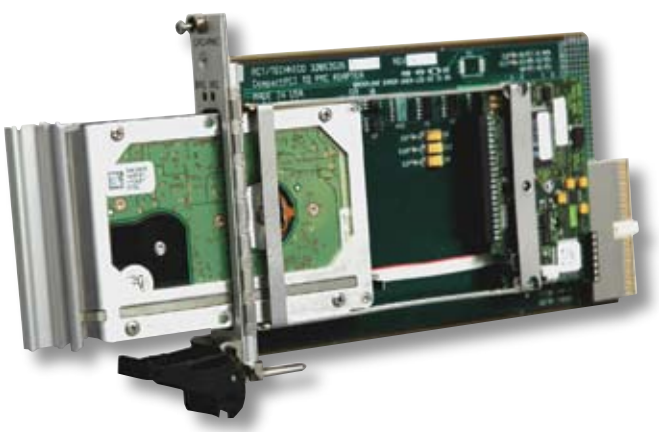

Figure 4 Removable hot swap disk shuttle.

Lastly, the good old command line interface can be used via Telnet, or the user can make use of e-mail facilities where a mail server can send the event or status information to an administrator. The end use environment will typically determine the management interface required.

Tradeoffs in NAS performance will be apparent when comparing network attached storage with direct attached storage. Ethernet-based systems must take into account the overhead of the transport protocols used to access the storage device. Performance will be a function of the disk performance, the processing power of the controller and the speed of the network interface. A more significant factor affecting performance is the number of times the data is duplicated. Using Gigabit Ethernet interfaces and 7,200 RPM drives results in performance as shown in Table 1.

Automatic data duplication and RAID services present additional overhead, but in system applications where moderate bandwidth is necessary, these services can provide the data availability needed for critical applications. In addition to bandwidth, it is necessary to consider storage availability and uptime as well as failover recovery metrics. Typical system interrupt times, including recovery from the failure of a drive, recovery from failure of a blade and recovery from the failure of a network interface, are outlined in Table 2.

Component Reliability and Environmental Considerations

At the core of the NAS solution is the reliability of the system components. Basic considerations include the construction and grade of the components used, including the rotating disk drives. In order to provide highly available storage, the storage devices need to be rated for continuous operation, or 24/7 operation, implying that the drive is powered on and rotating continuously.

The drives must also support a higher access rate. Drives designed for continuous use are designed for 50% duty cycles, meaning they are read or written to 50% of the time. Typical 2 ½" ATA drives are rated for a 20% duty cycle and are not for continuous operation. A typical drive's life under continuous operation would be 30,000 Power on Hours, or about four years, and MTBF would be in excess of 500,000 hours per drive. Typical blade reliability would exceed 150,000 hours per blade. The hardware is configured to provide a better than five nines storage solution. However, the actual availability must be considered at the system level, taking into account the network interface and operating software. Practical overall availability would be four nines or better.

Blade-level NAS implementations provide an advantage over conventional NAS solutions when considering the operating environment. The blades themselves are resistant to vibration as well as shock, and the drives can perform operationally at 10g,

Failed Component	System Interrupt Times	Method
Hard Drive	Instantaneous	Occurs automatically
Blade: Primary to Secondary blade	10 to 15 seconds	Occurs automatically
Blade: Secondary to Primary blade	15 to 20 seconds	Occurs automatically
Network (switch)	Less than 2 seconds	Occurs automatically

Table 2 Typical System Interrupt Times

11 ms half-sine. The operating temperature for rotating media is from 5° to 55°C (40° to 130°F), which meets the NEBS level 3 short-term temperature excursion requirement. Hence, board-level implementation benefits from the robust nature of the form-factor. In addition, blade-level NAS implemented via a PICMG 2.16 / VITA 31.1 packet switching backplane eliminates failure-prone cabling.

Blade-level NAS provides a scalable storage building block for embedded NAS servers. The NAS blade is ideal for transaction-based applications where data availability is essential. Centralized access to user data from network nodes is achieved using the client/server model that simplifies management and access of the user data. With an appliance-like interface, NAS at the blade level provides a compact, maintainable storage approach for embedded system applications. ◢

ACT/Technico
Ivyland, PA.
(215) 957-9071.
[www.acttechnico.com].

Get your next embedded application to the *finish line* faster...

Don't learn about them *on the street*

***Drive* your knowledge to the next level**

Get trained by the experts at

MCC Systems

	Boston, MA	Los Angeles, CA
Developing for Embedded Linux®	Sep 20-23 Nov 1-4	Oct 18-21 Jan 10-13
VxWorks® 5.x and Tornado® 2.x	Oct 18-21	*call for date*
Developing Embedded Linux Device Drivers	Sep 27-30	Oct 25-28
Eclipse™ Plug-in Development	Oct 11-14 Dec 6-9	Nov 8-11 Jan 17-20
VxWorks BSPs and Device Drivers	Jan 10-13	*call for date*

For complete course descriptions go to www.mccengineering.com/traininglist.htm
To register or find out about custom training call 940-206-5475 or email training@mccengineering.com

High-Availability Data Management for Real-time Applications

The performance and predictability needs of embedded systems often dictate the use of in-memory database systems. But for applications that must be highly available, how can a memory-based database offer the ability to survive the failure of its hardware or software environment?

by Steven T. Graves
McObject

High availability for in-memory data management is usually achieved through redundancy, or maintaining two or more synchronized databases in an active/standby configuration. In case the original primary database fails, continued availability is maintained by causing one of the standby instances to "take over" as the active database. A high-availability data management configuration is classified as either a one-safe or two-safe system depending on the approach used to propagate updates from the primary to the standby replica database(s).

The two-safe replication technique is based on synchronous communication between an active database and the standby database(s), and a voting protocol to determine whether the communication is successful. The communication represents a transaction, or a group of database read or write operations, which must succeed or fail as a group (Figure 1). The vote determines whether the transaction is applied (committed) or rolled back (aborted). Voting itself can be strong or weak. In strong voting there is mutual "agreement" between the database master and replicas that the transaction has been committed or rolled back. In weak voting, the active or "master" database decides unilaterally.

One-safe replication employs asynchronous communication between an active database and the standby database(s). The replication takes place outside the scope of the database transaction, and changes can commit on one database without necessarily succeeding on all others. Therefore, voting is by definition unilateral (Figure 2). One-safe and two-safe replication represent trade-offs between safety and availability.

Two-Safe Replication

Two-safe, or "eager" replication provides the highest degree of safety, since data is replicated within the scope of a transaction and there is no risk that once the transaction is committed to one database copy, it will fail to be propagated and replicated on the others.

Two-safe replication enables the application to distribute its query processing load across all available database copies. Because each transaction is committed to the entire group of databases, and every database copy in the group has the most recent data, a query issued against any copy will return the same results. Query performance can be improved if one of these

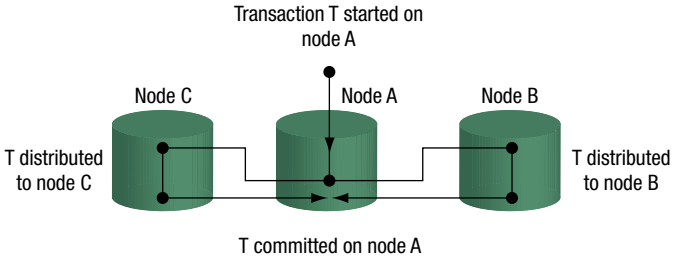

Figure 1 Two-safe, or "eager" replication has the advantages of updating as soon as a transaction occurs, high transaction durability and strong data consistence. However, it suffers from degraded transaction throughput.

Rolling window provides the flexibility to force quiescence only when the quantity of changes falls below a specified threshold, enabling the developer to minimize the quiescent period and maximize availability.

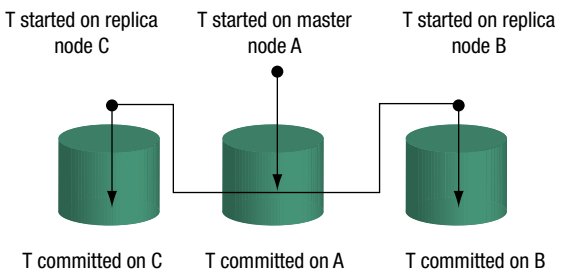

Figure 2 One-safe, or "lazy" replication updates asynchronously with less durability and weaker consistency, but offers higher transaction throughput.

database copies is local to the application, so that the application avoids communicating with the active or "master" database across the network.

Transactions are committed to the group of databases together or not at all. So if the system hosting the master database dies and the system needs to failover to a replica, the process can execute immediately. There is no need to wait for asynchronous transactions to "catch up," as could be the case with one-safe replication.

In two-safe replication, database transactions are completed (committed) in a process that requires synchronized message-passing among multiple databases. This requires longer resource holding time than one-safe replication. With two-safe replication, locks on the active database cannot be relinquished until the transaction commit message has been delivered to and (in the case of mutual voting) acknowledged by all involved database nodes.

This kind of message-passing incurs network latency, which inserts unpredictability into the transaction execution path. This can be a serious issue, because many real-time applications absolutely require predictable performance. A solution that meshes the greater safety of two-safe replication with the need for predictability is to build time-cognizance into the transaction commit protocol so that the application can assign time constraints to each phase of communication. In this way, predictability is ensured (because a process waiting on a message will not wait longer than the specified time).

One-Safe Replication

One-safe, or "lazy" replication provides the highest degree of availability, with reduced safety. Availability is maximized because the active database is not required to hold the resources (in database parlance, "locking") after committing the transaction itself. Therefore, the database resources are available for the next operation sooner, by virtue of not having to wait for network communication roundtrips. Depending on system requirements, a one-safe implementation may still be able to distribute its query processing load among the database group members, but this risks reading stale data, since some replica databases may not yet have been updated.

One-safe replication gains availability at the expense of safety. Since transactions are committed to individual databases in the group separately rather than together, a transaction that is committed to the active database may not be successfully propagated to every—or any—standby database(s). When an application relies on a purely in-memory database, this can happen if the system that hosts the active database fails after the transaction is committed but before it is transmitted to the standby database(s). Transaction logging can mitigate this risk by writing changes introduced by the transaction to persistent media such as a hard disk, flash or non-volatile RAM. One-safe (asynchronous) replication can cause a slight delay in the failover to a replica if the failover procedure has to wait for queued-up transactions to complete before converting the replica into the new master database.

The Right "High Availability" to Suit the Application

Apart from the distinctions between one-safe or two-safe replication, the term *high availability* has different interpretations for different types of applications. For a network infrastructure device such as a router, high availability means a percentage of time—commonly 99.999% or "five nines"—that the application must be available. For airborne or mission-critical military and defense systems, high availability means an *inability to fail* One or more redundant systems must remain available to continue processing if the primary system fails.

For data management, these different interpretations lead to different approaches to a major system capability: maintaining high availability when new replicas must be attached and synchronized.

When a new replica attaches to an active database, the active database must provision the replica with its current state.

This initial "snapshot" of the database must represent a transaction-consistent state, i.e., it cannot contain partial (uncommitted) transaction changes. The simplest way to achieve this is to defer the start of new transactions on the active database until the initial synchronization of the replica is complete. This is called "cold" initial synchronization.

To accomplish this, the database management system must lock the active database, or impose a quiescent period. However, blocking the primary database during this period reduces data availability, which defeats the purpose of maintaining redundant databases for high availability. In contrast to this static approach, the more sophisticated "hot synchronization" provisions the newly attached replica without interrupting normal transaction processing on the active database.

For many military and aerospace systems, cold synchronization may be satisfactory because such applications do not require "hot-swapping" of components that include embedded replica databases. Generally speaking, the device or craft is made mission-ready by bringing all the systems online and provisioning them with all necessary mission data. At this time, cold synchronization is sufficient. During the mission itself, no new replica databases are brought online.

In contrast, network infrastructure is frequently reconfigured during operation. For example, such devices often accommodate hot-swappable boards or blades. A router might use a chassis to hold several blades running different protocols, and if one board fails, a surviving board can "fill in" for it until the original is replaced. In another scenario, a new board might be installed to add an extra CPU to the router, or a desired new capability.

Installation must occur without shutting down the system, but the new board needs to be provisioned with the current data by an initial synchronization process. To uphold the five nines or higher availability mandate, an active database running within the router's embedded software cannot afford to defer new transactions. This and similar scenarios require hot initial synchronization.

In cold synchronization the active database is locked for the period required to transmit the current data to the replica(s). Every request to modify the database waits in a queue until the synchronization completes (when the active database is notified by the replica that it can continue accepting new data). The database may still be available for read-only access. For obvious reasons, the database management system (DBMS) must seek to minimize the locked time.

Hot synchronization, sometimes known as "rolling window," involves setting a "hot synchronization mode" and moving through the database organization units (rows, pages, files, etc. depending on the database implementation) and transferring these units to the replica. Normal application processing that "touches" the database causes the database run-time to set a "dirty" flag on an organization unit when its contents have been modified.

If the dirty flag is set on a unit that the hot synchronization process has not reached yet, it is of no consequence; the unit would have been transferred in its changed state anyway. When the hot synchronization process reaches the end of the database it wraps around to the beginning and looks for units that have the dirty flag set. This indicates that the unit was modified after the hot synchronization process transferred that unit to the replica, so it must be transferred again, with the new changes.

Obviously, if new dirty flags keep arriving, this process could endlessly loop and initial synchronization might never complete. The application might at some point need to force quiescence in order to complete the initial synchronization. However, thresholds can be set, and quiescence imposed only when the quantity of modified, un-transferred units falls below a given value. Alternatively, quiescence can be imposed after a given amount of time has elapsed. The duration of the period of forced quiescence can be kept to a minimum, and rates of change and network transport speeds can be considered in the determination of appropriate thresholds.

Rolling window provides the flexibility to force quiescence only when the quantity of changes falls below a specified threshold, enabling the developer to minimize the quiescent period and maximize availability. This method is architecturally uncomplicated, results in small code size and is less error-prone. As with the choice between one-safe vs. two-safe replication, these differing approaches to initial synchronization give the system designer some flexibility in implementing a highly available database system, enabling productive tradeoffs by developers who are willing to prioritize their goals. ◢

McObject
Issaquah, WA.
(425) 831-5964.
[www.mcobject.com].

524 MHz, Low-Power Multimedia ARM Computer on a PC/104

Combining a flat panel display interface with a fast, low-power ARM processor in a PC/104 form-factor, the SBC1670 from Micro/sys will operate from -40° to +85°C. The CPU has multimedia capabilities, which include support for an 800 x 600 color flat-panel display, audio output and de-bounced keypad input. The Intel PXA270 processor is based on the ARM5 core and implements a super-pipelined RISC architecture while adding a number of integrated peripherals. The PXA270 can run at up to 524 MHz and can dynamically change speed in response to performance or power consumption needs. On-chip cache, a watchdog timer, an SDRAM controller, a CompactFlash interface and a USB host controller are also integrated on the same silicon.

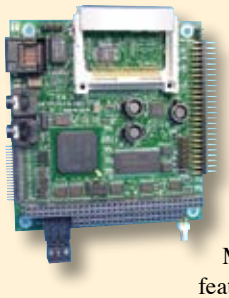

With 128 Mbytes of SDRAM, and a 64 Mbyte resident flash array, the SBC1670 also features five serial ports and a 10/100BASE-T Ethernet controller to handle a variety of communication needs. Four of the serial ports have RS232 transceivers, while the remaining serial port is configured for RS485 communication. The SBC1670 can boot Linux, Windows CE and VxWorks from its onboard flash. If more I/O is needed, the SBC1670 allows expansion through its CompactFlash socket, which supports storage devices and I/O devices, such as Wi-Fi cards. A free development kit is provided that includes cables, sample software and full documentation.

The basic SBC1670 starts at $495 in single quantity. An industrial temperature version starts at $575.

Micro/sys, Montrose, CA. (818) 244-4600. [www.embeddedsys.com].

Rugged, Embedded PC Delivers High Performance in Harsh Conditions

Remote monitoring, industrial automation and communications systems often need a rugged PC that can deliver high performance. GE Fanuc's Talon 8400, a self-contained Intel Celeron-based PC, provides high-performance computing in a small space. It operates in harsh conditions at 5° to 40°C, which can be optionally extended to a range of -20° to 50°C.

The Talon 8400 features 256 Mbytes of PC100/PC133 SDRAM and a 20 Gbyte hard drive, and has certifications for UL, CE, NEBS level 3 and FCC B. Support for Windows XP Professional is available. For connectivity and expansion, the PC includes two fast Ethernet ports, four USB 2.0 ports, an SVGA connector, an IEEE 1284 port, four individually isolated RS232/485 serial ports and two open PCMCIA slots.

The Talon 8400 can be DIN rail, panel or tabletop mounted, and used as a stand-alone computer or embedded in a larger system. A UL 1604 option is available for hazardous locations, in addition to standard features such as fanless operation and electronic shutdown with an over-temperature LED for system overheating. Pricing ranges from $2,000 to $2,600.

GE Fanuc Embedded Systems, Huntsville, AL. (800) 322-3616. [www.gefanuc.com].

Pentium 4 Industrial Motherboard is Intel VRD 10.1-Compliant

Demanding, industrial-grade applications such as factory automation and computer automation systems and medical devices must be upgraded with new processors. With that in mind, the ATX-865G P4 Industrial Motherboard from Arista is Intel VRD 10.1-compliant to support future processors.

The ATX-807 CPU includes a 400/533/800 MHz FSB Pentium 4 Prescott (90 nm) CPU and four DDR DIMM sockets that support up to 4 Gbytes of DDR 266/333/400 unregistered non-ECC memory. The board features an 865G integrated Intel Extreme Graphic Engine with a 266 MHz core frequency and onboard 100/10 Intel 82547GI and 82541GI Ethernet controllers.

The ATX-865G board has eight USB 2.0 ports that support USB hot-plug. For expansion, there are one high-speed parallel port with SPP/EPP/ECP mode support, four 16550 UART-compatible ports, a pin-header connector for an optional IrDA external connector, two Enhanced PIDE interfaces and support for up to two SATA-150 HDDs. An FDD interface supports up to two floppy drives, and two ISA slots are provided for applications that use ISA legacy cards. Pricing starts at $365.

Arista, Camarillo, CA. (510) 266-1800. [www.aristaipc.com].

Rugged ADC PMC Module Targets SDR Challenges

ICS, a part of Radstone, has announced its new ICS-8554 ADC PMC module. This comes as a response to demand for analog to digital conversion (ADC) solutions that not only offer performance and flexibility, but are rugged enough to deploy in the harshest environments, and yet are available as COTS products. The module is designed for intermediate frequency (IF) Software Defined Radio (SDR) applications such as multiple element receive beamforming, coherent radar or secure communications. The ICS-8554 combines DSP and ADC technologies to bring a small form-factor, lightweight—capable of vehicle or man pack mounting—solution.

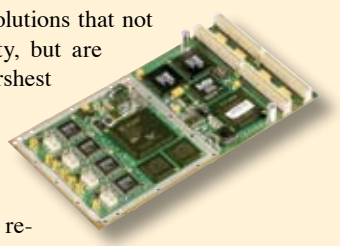

At the heart of the ICS-8554 is the Analog Devices AD6645 high-performance bipolar ADC, which is capable of sustaining its 80 MHz sample rate at temperatures as low as -40ºC and as high as +85ºC. An onboard 3M gate Xilinx Virtex II FPGA (XC2V3000) provides user programmability as well as enabling processing closer to the antenna, resulting in higher system throughput. The 8554 features a 70 dB signal-to-noise-ratio and 80 dB spur-free dynamic range (SFDR), for maximum signal purity and integrity. Pricing of the ICS-8554 begins at $8,500.

Radstone Embedded Computing, Ottawa, Ont. (613) 749-9241. [www.radstone.com].

Ultra-Rugged Portable PXI Test Set Targets Field Testing Apps

Engineers in the field often need to perform PXI-based data acquisition and testing under harsh conditions. The Geotest MTS-207 is an ultra-rugged, portable PXI-based configurable platform for field testing and data acquisition systems.

Inside its rugged enclosure, a 14-slot PXI chassis mainframe has one slot dedicated to a 6U PXI controller. Slots two through seven can accommodate either 3U or 6U PXI/CompactPCI instruments. Slots eight through fourteen accommodate only 3U PXI/cPCI instruments. Available controllers for the MTS-207 include 1.2 GHz and 1.7 GHz Pentium 4 CPUs. The MTS-207 also provides sufficient space for custom circuitry, such as specialty power supplies and circuit cards. Fan assemblies provide dedicated cooling to the PXI chassis and internal circuitry. Multiple shock mounts and extensive internal temperature sensors ensure reliable operation under adverse operating conditions.

An optional Remote Control & Display Unit (RCDU) combines an LCD display with a touch panel for operation from a distance of up to 25 feet in environments where a keyboard and mouse are not functional or practical. Optional built-in heaters for the PXI chassis and the RCDU allow the MTS-207 to operate in extremely low temperatures. Prices begin at $25,000.

Geotest, Irvine, CA. (949) 263-2222. [www.geotestinc.com].

PC/104 CPU Offers Onboard Memory and Data Acquisition

A single-board PC/104 CPU integrates a Pentium-II class CPU running at 200 MHz with 128 Mbyte memory, Ethernet and a professional-quality analog and digital I/O circuit with auto-calibration. The Elektra from Diamond Systems supports Linux, DOS, QNX and Windows CE operating systems and consumes only 5.5 watts to operate over the -40° to +85ºC extended temperature range without a fan.

Elektra's built-in data acquisition circuit provides stable A/D with 16-bit resolution and programmable input ranges. A 512-sample FIFO and onboard programmable timer enable an error-free 100 KHz A/D sampling rate in any operating system. It includes four 12-bit D/A channels, 24 programmable digital I/O lines and two programmable counter/timers. All data acquisition features are fully supported by Diamond's Universal Driver software. The built-in auto-calibration circuitry provides enhanced accuracy for analog measurements.

Two configurations are available: one includes CPU plus data acquisition, while the other is the lower-cost CPU-only configuration. Both versions include four RS-232 ports, two USB 1.1 ports with USB floppy support, 10/100 Mbit/s Ethernet, PS/2 keyboard/mouse, parallel and IDE ports, programmable watchdog timer and onboard backup battery for real-time clock and BIOS settings.

In quantity 1-9, the ELK200-EA-XT with data acquisition is priced at $750, while the ELK200-E-XT is $450.

Diamond Systems, Newark, CA. (510) 456-7800 [www.diamondsystems.com].

Motion Control Card Provides Ultra-Fast Servo Loop Rate

High-performance motion control applications, such as medical, scientific and high-performance automation, need especially fast servo loop rates of under 100 microseconds. To fill that need, Performance Motion Devices has released the Magellan-PCI Motion Controller card, with a faster servo loop rate of 50 microseconds per axis.

The Magellan-PCI Motion Controller is available in 1-, 2-, 3- and 4-axis versions. It supports DC brush, brushless DC, microstepping and pulse and direction motors. In addition to a servo loop rate of 50 microseconds per axis, Magellan-PCI provides a 5 Mcount/sec quadrature encoder input rate, and pulse and direction output of up to 5 Mpulses/sec. Other features include trajectory generation, servo loop closure, quadrature signal input, motor output signal generation, servo trace, commutation and the ability to make on-the-fly changes. Communication occurs through a PCI bus, via a serial port or using CANBus.

Designers can use the card to create cost-effective, high-performance motion systems using standard C or C++ programming. The card accepts position, velocity, acceleration, deceleration and jerk input parameters and automatically generates the programmed trajectory. On-the-fly changes and PLC-style external signal input/output breakpoints can be used to create application-specific profile changes. Prices start at $638 in OEM quantities.

Performance Motion Devices, Lincoln, MA. (781) 674-9860. [www.pmdcorp.com].

Universal Network Analysis Tool for Controller Area Networks

Designers of systems that use CAN networks face the demands of growing system complexity. To help alleviate this situation, Accurate Technologies has developed CANLab, a universal network analysis tool for monitoring and logging messages and signals generated on a CAN.

Designed using .NET technology, CANLab provides the ability to send messages over any available network channel, facilitating the ability to re-send recorded bus traffic. Additionally, real-time data and data imported from previous recordings can be viewed simultaneously for in-depth post-analysis. Messages and signals transmitted from any available network channel can be displayed and manipulated in a variety of ways. The CANLab Workspace screens can be configured through a docking window management scheme.

The tool offers a full-featured scripting engine, letting users programmatically interact with the CANLab run-time environment. Complex functions can be composed to respond to several different types of events. CANLab supports Windows 2000 and XP, is compatible with C# and requires a processor running at 500 MHz or faster with at least 128 Mbytes of RAM. It supports a wide variety of CAN interface hardware, including Kvaser v3.8 and Vector v4.3, as well as both .DBC and .UEF database formats. Prices start at $2,000.

Accurate Technologies, Wixom, MI. (617) 739-8693. [www.accuratetechnologies.com].

Rugged PMC Modules Offer 3D Graphics Processing

A line of graphic PMCs based on ATI Technologies' Radeon Mobility 9000 (M9) mobile graphics processor supports 2D and 3D acceleration, OpenGL and DirectX. The rugged G2 line from SBS is targeted for military and avionics applications such as cockpit avionics and terrain overlay systems, or any application that requires a small footprint and high display performance. The Family consists of three models: the G2 Basic PMC, the G2 Dual PMC and the G2 Plus PMC.

The economical G2 Basic PMC is designed for applications that need only a single display output channel. The G2 Basic provides a wide range of display output options including DVI with up to 1600 x 1200 pixel resolution, LVDS, RGB, NTSC, PAL and S-Video. The G2 Dual PMC features the same video display output support options as the G2 Basic PMC and adds another display channel for those applications that require dual displays. The G2 Plus PMC provides the broadest range of capabilities. Dual display outputs plus up to two channels of video input capture make the G2 Plus PMC an ideal choice for demanding video capture applications. The G2 Plus PMC's addition of graphics overlay capabilities is especially useful for applications requiring display and manipulation of multiple video streams. The modules are supported by optional OpenGL drivers from ALT Software Inc., which are available for a number of different operating systems including Linux, VxWorks, Integrity and Windows XP.

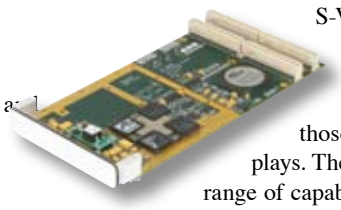

Pricing for the G2 Graphics PMC Module in single quantities starts at $2,200.

SBS Technologies, Albuquerque, NM. (801) 483-1533. [www.sbs.com].

1.3 GHz EBX SBC Comes in at Lower Cost

The new Celeron M version of the Cobra SBC from VersaLogic offers 1.3 GHz processing at a cost about 25% lower than the 1.6 GHz model. The Celeron M Cobra is engineered for all applications where performance and reliability are critical, including aerospace, medical, government, security and transportation. The Celeron M Cobra includes integrated Extreme Graphics 2 video with flat panel support, dual 10/100 Ethernet, up to 2 Gbyte DDR RAM, audio, USB 2.0, 4 COM ports, 32 lines of digital I/O, LPT interface, IDE interface and a CompactFlash socket. Additional expansion is easily accommodated via ISA and PCI busses on the PC/104-Plus expansion site.

The Cobra Celeron M single board computer is manufactured to quality and reliability standards that include safety-critical features such as TVS devices for ESD protection, VCC-sensing reset circuit and a watchdog timer for hardware-level application control. The Cobra EBX-12g is currently available from stock and includes VersaLogic's 5-year availability guarantee. Pricing is under $1,200 in low OEM quantities.

VersaLogic, Eugene, OR. (541) 485-8575. [www.versalogic.com].

EPIC SBS Targets Machine-to-Machine Applications

Specifically targeting machine-to-machine (M2M) applications that require a processor, a wide variety of I/O and connectivity, the EPX-GX from WinSystems supports eleven onboard I/O options plus more available from PC/104 modules that can be added on. The EPX-GX is an AMD Geode 500@1.0W-based, EPIC-compatible single board computer (SBC). AMD Geode processors have extremely low power dissipation, which allows fanless operation from -40° to +85°C. Applications include test equipment, medical instruments, communications devices, transportation systems, military/COTS, data loggers, security, robotics, semiconductor manufacturing instruments and industrial control systems.

The board is configured with up to 512 Mbytes of PC2700 DDR SDRAM plus a CompactFlash memory socket. It also has a 10/100 Ethernet controller, video with CRT and flat panel interfaces, four serial COM channels, 24 digital I/O lines, 6 channels of AC97 audio, and the standard AT peripheral feature set is included. There is a socket for bootable Type I and II CompactFlash cards. The EPX-GX2 also includes a socket for a miniPCI wireless 802.11 a/b/g card. A connector is included to support a remote GPS receiver. The EPX-GX runs both 16-bit and 32-bit x86 instruction set software. Pricing is $499 in OEM quantities.

WinSystems, Arlington, TX. (817) 274-7553. [www.winsystems.com].

Tiny Embedded Processor/Device Server Bridges M2M and Datacom Networks

To bring powerful, data-center grade information exchange capabilities to embedded devices arrayed throughout private networks and the Internet, engineers must be able to quickly and easily build Ethernet or IP networking capabilities into their products. The XPort AR processor/device server module from Lantronix delivers this ability in a miniaturized RJ45 package.

The XPort AR module is the company's first product to include its enterprise-class Evolution Network Operating System (ENOS). It includes a compact, hardened, RTOS kernel with a full Layer 3 TCP/IP networking stack and a CGI-based dynamic Web server, along with the Lantronix DSTni 120 MHz processor, an Ethernet 10/100 interface, up to three serial ports, 11 general-purpose I/O control pins, 4 Mbytes of flash memory and 1.25 Mbytes of SRAM. The module includes an extensive suite of enterprise-grade, open standards-based security such as SSL 3.0 and SSH2 encryption protocols. The XPort AR also provides XML and RSS feeds for configuration and information transport.

An evaluation kit includes the module, an evaluation board, a universal power supply with snap-fit plugs for different countries, CAT5 network cable, a serial adapter, an RS-232 cable and a complete user manual. Also included is a CD with sample code, application notes, CAD PCB files and software utilities. Prices start at $70 in single quantities.

Lantronix, Irvine, CA. (800) 526-8764. [www.lantronix.com].

PowerPC and Microblaze Development Kit for Xilinx Virtex-4 FPGAs

The PowerPC and MicroBlaze Development Kit, Virtex-4 FX12 Edition is a design environment that enables embedded developers to create processor-based systems on Xilinx Virtex-4 FPGAs. The kit delivers an integrated platform with hardware, design tools, intellectual property (IP) and reference designs. The Development Kit supports both the PowerPC 405 hard processor and MicroBlaze soft processor for the Xilinx Virtex-4 FX series of Platform FPGAs. The kit integrates a Virtex-4 ML403 development board (hardware) and the Platform Studio embedded tool suite and Integrated Software Environment FPGA design software. It additionally provides more than 60 IP cores and a JTAG probe along with serial and Ethernet cables, pre-configured flash device and a variety of pre-verified reference designs. Central to the functionality is the Virtex-4 FX12 FPGA device with the immersed PowerPC 405 processor and integrated 10/100/1000 Mbit/s Ethernet controllers.

The programmable platform and Platform Studio design enable developers to craft embedded systems and optimize the combination of features, performance, area and cost. Developers can choose the most effective processor core for the target application, customize IP, optimize the performance and validate software on a development board before their own custom hardware is even back from the shop. In addition, the pre-verified reference designs offer a set of options for rapidly configuring complete systems. These include both hardware and software implementations to exercise features of the ML403 platform, such as Ethernet, DDR memory, video and audio functions, as well as Wind River Systems' VxWorks and Monta Vista embedded Linux operating system demonstrations. The kit is priced at $895.

Xilinx, San Jose, CA. (408) 559-7778. [www.xilinx.com.].

PC/104-Plus 4-Port USB 2.0 Controller Supports uDiskonChip SSD

A PC/104-Plus SBC from Parvus is a 4-port, high-speed USB 2.0 host controller board capable of directly powering attached USB devices. The COM-1440 now also comes with onboard mounting support for two M-Systems uDiskOnChip (uDOC) Solid State Disks (SSD). Delivering USB data transfer speeds of up to 480 Mbits/s, the COM-1440 features four downstream USB ports, along with two onboard SSD mezzanine sockets for up to 4 Gbytes of secure high-speed flash memory. Notably, two of the product's USB ports are enhanced to supply two amps of drive for directly powering attached USB mass storage and other devices through the USB cable.

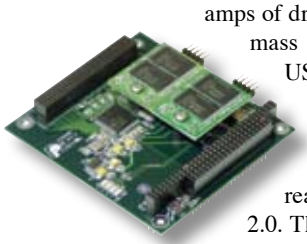

The new uDOC-compatible sockets enable rugged solid state storage with throughput rates of 20 Mbyte/s read and 10 Mbyte/s write using USB 2.0. The two memory device sockets are connected in parallel with two of the external USB ports. When the two uDOC sockets are populated, one standard (500mA) USB port and one enhanced drive USB port are free for system expansion. Accessible as independent logical disks and available in extended temperature formats, uDOC encrypted solid state flash disk devices from M-Systems are available in capacities ranging from 64 to 2048 Mbytes each. Pricing is $132 each in 100-unit quantities.

Parvus, Salt Lake City, UT (801) 483-1533. [www.parvus.com].

Rugged CompactPCI SBC with MPC5200 PPC Features FPGA

Embedded mobile or industrial control applications, such as transportation systems, require lots of functionality in demanding environments. With those needs in mind, Men Micro has introduced the F12, a rugged, single-slot 3U CompactPCI SBC based on a 384 MHz MPC5200 PowerPC. The F12 includes an Altera Cyclone FPGA for graphics or other custom I/O.

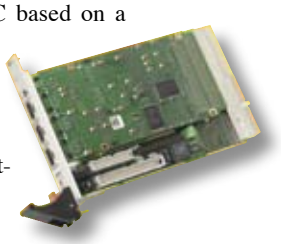

The MPC5200 processor includes a telematics communications unit, floating point unit, memory management unit, DRAM controller and the BestComm/DMA I/O controller, which can drive a number of industrial interfaces such as CAN, USB, Fast Ethernet and SPI. Onboard storage resources include up to 256 Mbytes of SDRAM, 1 Gbyte of NAND flash memory, 2 Mbytes of static RAM and 16 Mbytes of graphics RAM. For expansion, two Fast Ethernet ports, a serial communication (COM) port and one USB interface are included. An optional version has robust D-Sub connectors in lieu of the standard RJ45 connectors. Two CAN channels are also available with the addition of Men Micro's small SA-Adaptor transition modules.

The F12 operates over the extended industrial temperature range of -40° to 85°C. To help reduce the effects of shock and vibration, all onboard memory is soldered to the PCB. Comprehensive board support packages based on Men Micro's BIOS for PowerPC processors, MENMON, are available for Linux, VxWorks and QNX. Pricing starts at $954 for single units.

Men Micro, Dallas, TX. (512) 267-8883. [www.menmicro.com].

Channel-to-Channel Isolated Data Acquisition Module for Industrial Applications

Industrial applications requiring highly accurate, small-signal measurements with high common-mode voltages need measurement systems that can be easily protected from possible voltage spikes or power surges. The NI PXI-4224 channel-to-channel isolated data acquisition module from National Instruments was designed with these needs in mind.

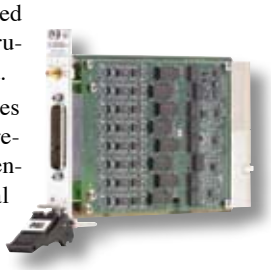

The 60 VDC, Cat. I NI PXI-4224 provides eight channels of 16-bit analog input measurements. Channel-to-channel isolation lets engineers take measurements in an industrial setting where voltage spikes and transient signal inputs often occur. This isolation also prevents ground potential loops and provides improved noise immunity. Designed to work with LabVIEW 7 Express and NI-DAQmx measurement services software, the module is compatible with Visual Basic, C/C++ and C#, and supports Windows 2000/NT/XP and LabVIEW Real-Time.

The NI PXI-4224 features integrated signal conditioning and provides programmable gain settings per channel to ensure maximum measurement accuracy. A 333 kS/s sampling rate and 16-bit resolution offer the flexibility to measure a broad range of signals with highly accurate voltage measurements. Other features include direct sensor/signal connectivity, a NI-PGIA custom instrumentation amplifier for accurate measurements, temperature drift protection circuitry, a PXI trigger bus to synchronize operations between two or more modules and digital triggering. Pricing begins at $1,995.

National Instruments, Austin, TX. (800) 813 3693. [www.ni.com].

Advertiser Index

Company	Page	Website
Acromag	57,59,61	www.acromag.com
ACT/Technico	22	www.acttechnico.com
ADLINK Technology, Inc.	13	www.adlinktechnology.com.
American Arium	49	www.arium.com
APW Electronic Solutions	7	www.apw.com
Arcom Control Systems LTD.	60	www.arcom.com
BittWare	45	www.bittware.com
CDV Inc	58	www.cdvinc.com
Diversified Technology	23	www.dtims.com
Dynatem, Inc.	8	www.dynatem.com
Elma Bustronic Corp.	28	www.elmabustronic.com
ELMA Electronic, Inc.	10	www.elma.com
Embedded Planet	24	www.embeddedplanet.com
Enea Embedded Technology	38,39	www.ose.com
Interactive Circuits and Systems	75	www.ics-ltd.com
Interface Concept	6	www.interfaceconcept.com
Kontron America	31	www.kontron.com
MCC Systems	66	www.mccengineerming.com/traininglist.htm
Mercury Computer	65	www.mercury.com
Micro Memory LLC	76	www.micromemory.com
Micro/sys, Inc.	40	www.embeddedsys.com
One Stop Systems	54	www.onestopsystems.com
Phoenix International	6	www.phenxint.com
Real-Time & Embedded Computing Conf.	56	www.rtecc.com
Red Rock Technologies, Inc.	69	www.redrocktech.com
SBE, Inc.	37	www.sbei.com
SBS Technologies	2,4	www.sbs.com
Sealevel Systems	35	www.sealevel.com
Technobox, Inc.	44	www.technobox.com
Thales Computers	19	www.thalescomputers.com
Themis Computer	3,29	www.themis.com
VadaTech	18	www.vadatech.com
VersaLogic Corporation	51	www.versalogic.com
VME Product Showcase	32	
WinSystems	21	www.winsystems.com

RTC (Issn#1092-1524) magazine is published monthly at 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.. Periodical postage paid at San Clemente and at additional mailing offices.
POSTMASTER: Send address changes to RTC, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

Ultimate Flexibility. Ultimate Precision.

No wonder we're an Ultimate Product.

The performance, functionality and value for money of ICS products are world-beating. But you don't just have to take our word for it. The 1,000+ engineers who took part in the EE Times Ultimate Products program chose the ICS-645C as a finalist. Its high-speed, high-precision, and high-density differential data acquisition board gained the highest praise.

Designed for PCI applications in high-frequency sonar and high-speed test and measurement, its onboard front end signal conditioning capability allows simpler, more cost-effective and more compact solutions to be developed. Which all means that the real winners are ICS customers.

The Ultimate Generation is here! ICS-645C.

ICS
SENSOR PROCESSING

EETIMES' ULTIMATE PRODUCTS

ICS-645C chosen as a finalist
in the Ultimate Products supplement
of EE Times Feb 2005

CHECK THE FULL SPEC NOW AT
www.ics-ltd.com/645c

OR CALL TOLL FREE (US ONLY)
800-267-9794 or 613-749-9241

CoSine™ SoC

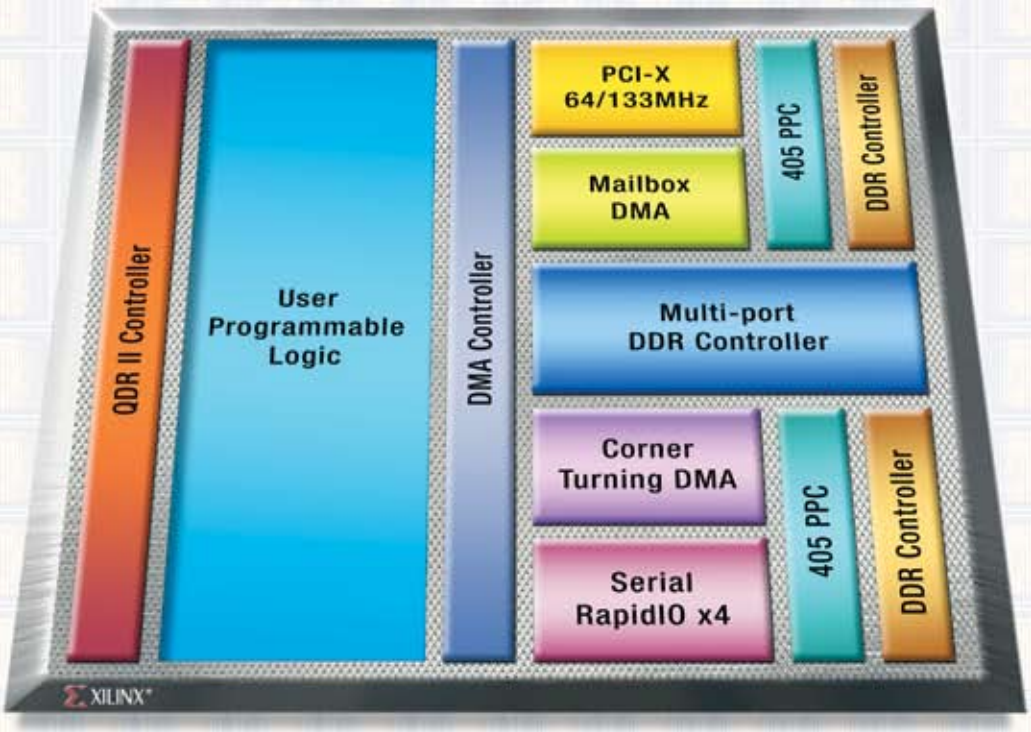

System-on-Chip for FPGA Processing

CoSine is a fully integrated, completely preconfigured System-on-Chip for real time FPGA based signal processing. Bridging two high speed interfaces (Serial RapidIO®, PCI-X®, or PCI-Express®) through a multi-port DDR controller, the elegant design enables non-contentious access to a User Programmable Logic (UPL) block.

Similar to a structured ASIC, the IP cores, memory controllers, specialized DMA engines, embedded processors and surrounding logic are factory preconfigured and supplied as a fully tested system. This provides users the ability to focus on application specific state machine processing in the UPL block without concerning themselves with coding other modules, complicated SoC integration, or verification.

The CoSine development toolkit includes a standalone ATCA® board with up to 8GB of DDR that demonstrates continuous sustained transfers from a 64-bit/133MHz PCI-X PrPMC site at maximum bandwidth through CoSine to a Serial RapidIO x4 XMC site. Along with a complete *"How to"* Developer's Manual, the kit includes a demo program using an off-the-shelf Xilinx 1024 FFT, robust library of VHDL test benches, and extensive suite of PCI, sRIO, and standalone diagnostic "C" test code.

CoSine will also be offered on several forthcoming XMC, AMC, and Othello VxS® VITA 41 and 46/48 formats in both commercial and conduction cooled options.

micro memory

MICRO MEMORY, LLC.
9540 Vassar Avenue, Chatsworth, California 91311 U.S.A.
(US) Tel 818 998 0070 • (US) Fax 818 998 4459
www.micromemory.com • sales@micromemory.com